



**HAL**  
open science

# An Equivalence between Bayesian Priors and Penalties in Variational Inference

Pierre Wolinski, Guillaume Charpiat, Yann Ollivier

► **To cite this version:**

Pierre Wolinski, Guillaume Charpiat, Yann Ollivier. An Equivalence between Bayesian Priors and Penalties in Variational Inference. 2024. <hal-02466702v2>

**HAL Id: hal-02466702**

**<https://hal.science/hal-02466702v2>**

Preprint submitted on 24 Feb 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

# An Equivalence between Bayesian Priors and Penalties in Variational Inference

Pierre Wolinski\*    Guillaume Charpiat†    Yann Ollivier‡

## Abstract

In machine learning, it is common to optimize the parameters of a probabilistic model, modulated by an ad hoc regularization term that penalizes some values of the parameters. Regularization terms appear naturally in Variational Inference, a tractable way to approximate Bayesian posteriors: the loss to optimize contains a Kullback–Leibler divergence term between the approximate posterior and a Bayesian prior. We fully characterize the regularizers that can arise according to this procedure, and provide a systematic way to compute the prior corresponding to a given penalty. Such a characterization can be used to discover constraints over the penalty function, so that the overall procedure remains Bayesian.

## 1 Introduction

Adding a penalty term to a loss, in order to make the trained model fit some user-defined property, is very common in machine learning. For instance, penalties are used to improve generalization, encourage sparsity in a model, or prune neurons and reduce the rank of the tensors of weights in the case of neural networks. Therefore, usual penalties are user-defined and justified mostly empirically, and integrated to the loss as follows:

$$\mathcal{L}(\boldsymbol{\theta}) = \ell(\boldsymbol{\theta}) + r(\boldsymbol{\theta}),$$

with  $\boldsymbol{\theta}$  the vector of parameters in the model,  $\ell(\boldsymbol{\theta})$  the error term and  $r(\boldsymbol{\theta})$  the penalty.

From a Bayesian point of view, optimizing such a loss  $\mathcal{L}$  is equivalent to finding the Maximum A Posteriori (MAP) of the parameters  $\boldsymbol{\theta}$  given the training data and a prior  $\alpha(\boldsymbol{\theta}) \propto \exp(-r(\boldsymbol{\theta}))$ . Indeed, assuming that the loss  $\ell$  is a log-likelihood loss, namely,  $\ell(\boldsymbol{\theta}) = -\ln p_{\boldsymbol{\theta}}(\mathcal{D})$  with dataset  $\mathcal{D}$ , then minimizing  $\mathcal{L}$  is equivalent to minimizing  $\mathcal{L}_{\text{MAP}}(\boldsymbol{\theta}) = -\ln p_{\boldsymbol{\theta}}(\mathcal{D}) - \ln(\alpha(\boldsymbol{\theta}))$ . Thus, within the

---

\*Laboratoire de Mathématiques d’Orsay, Université Paris-Saclay, France, pierre.wolinski@normalesup.org

†Tau Team, LRI, Inria Saclay, Saclay, France

‡Meta AI Research, Paris, France

MAP framework, we can interpret the penalty term  $r$  as the influence of a prior  $\alpha$  (MacKay, 1992b).

However, the MAP approximates the Bayesian posterior very roughly, by taking its maximum. Variational Inference (VI) provides a *variational posterior* distribution rather than a single value, hopefully representing the Bayesian posterior much better. VI looks for the best posterior approximation within a *variational family*  $(\beta_{\mathbf{u}}(\boldsymbol{\theta}))_{\mathbf{u}}$  of approximate posteriors over  $\boldsymbol{\theta}$ , parameterized by a vector  $\mathbf{u}$ . So, the *variational parameters*  $\mathbf{u}$  are trained instead of  $\boldsymbol{\theta}$ . Typically, one can choose, for  $\boldsymbol{\theta} \in \mathbb{R}^N$ ,  $\boldsymbol{\mu} = (\mu_1, \dots, \mu_N)$  and  $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_N)$ :

$$\beta_{\mathbf{u}} = \beta_{\boldsymbol{\mu}, \boldsymbol{\sigma}} = \mathcal{N}(\mu_1, \sigma_1^2) \otimes \dots \otimes \mathcal{N}(\mu_N, \sigma_N^2).$$

In VI, the loss to be minimized over  $\mathbf{u}$  is:

$$\begin{aligned} \mathcal{L}_{\text{VI}}(\beta_{\mathbf{u}}) &= -\mathbb{E}_{\boldsymbol{\theta} \sim \beta_{\mathbf{u}}} \ln p_{\boldsymbol{\theta}}(\mathcal{D}) + \text{KL}(\beta_{\mathbf{u}} \parallel \alpha) \\ &= \underbrace{\mathbb{E}_{\boldsymbol{\theta} \sim \beta_{\mathbf{u}}} \ell(\boldsymbol{\theta})}_{\text{data fit term}} + \underbrace{\text{KL}(\beta_{\mathbf{u}} \parallel \alpha)}_{\text{penalty term}}, \end{aligned} \quad (1)$$

which is also an upper bound on the Bayesian negative log-likelihood of the data (Jordan et al., 1999). Here  $p_{\boldsymbol{\theta}}(\mathcal{D})$  is the likelihood of the full dataset  $\mathcal{D}$  given  $\boldsymbol{\theta}$ ,  $\ell(\boldsymbol{\theta}) = -\ln p_{\boldsymbol{\theta}}(\mathcal{D})$  is the log-likelihood loss, and  $\alpha$  is the Bayesian prior. The variational posteriors  $\beta_{\mathbf{u}}$  that minimize this loss will be concentrated around values of  $\boldsymbol{\theta}$  that assign high probability to the data, while not diverging too much from the prior  $\alpha$ . Thus, the KL-divergence term can be seen as a penalty  $r(\mathbf{u})$ .

**Are we still Bayesian when using a custom penalty?** We study whether the prior-penalty correspondence above for the MAP extends to the richer variational inference setup. We assume that we optimize a VI loss of the type:

$$\mathcal{L}(\beta_{\mathbf{u}}) = \mathbb{E}_{\boldsymbol{\theta} \sim \beta_{\mathbf{u}}} \ell(\boldsymbol{\theta}) + r(\mathbf{u}),$$

and consider the following question: given a variational family  $(\beta_{\mathbf{u}})_{\mathbf{u}}$  and a penalty  $r(\mathbf{u})$ , can  $r(\mathbf{u})$  be interpreted as a KL-divergence  $\text{KL}(\beta_{\mathbf{u}} \parallel \alpha)$  for some prior distribution  $\alpha$ ? We answer this question in Section 4, with Theorem 1 and Corollary 1, and we provide a formula for the corresponding prior  $\alpha$ , if it exists. With these theoretical results, it is possible to check whether a setup with a custom penalty  $r$  is actually Bayesian. Moreover, if it is the case, we provide a formula to compute the corresponding prior  $\alpha$  explicitly.

**How to tune the shape of a penalty to make it fit the Bayesian setup?**

Let us consider a variational family with a penalty containing several undetermined hyperparameters (possibly functions). Results of Section 4 can be used to find constraints over these hyperparameters to make the penalty equal to a KL-divergence with a prior  $\alpha$ . That is, we provide a hyperparameter tuning method based on Bayesian principles, and we use it in several examples in Section 5.

Our main contributions are five-fold:

1. we show how to interpret a custom penalty as the influence of a Bayesian prior, if possible, and we propose a formula to compute the prior (Theorem 1);
2. under some conditions over the variational family and the search space of the prior, we provide a set of necessary and sufficient conditions that the penalty must fulfill to ensure existence and uniqueness of a corresponding prior (Corollary 1);
3. we apply these results to a series of examples, showing that, among other things: if the penalty is  $\mathcal{L}^2$  and corresponds to a prior, then this prior is Gaussian (Corollary 4); if the penalty is  $\mathcal{L}^1$  and the candidate variational posteriors have a smooth density, then no corresponding prior exists (Proposition 2);
4. additionally, we provide a result about sums of penalties, when each one involves the mean of a different parameter: if this sum corresponds to a Bayesian prior, then this prior can be decomposed into a product of independent priors, each one over a different parameter (Corollary 2);
5. we show that these theoretical results can be used to tune the strength of the penalty when training neural networks in a VI setup.

**Outline.** Section 2, provides a short review of the related works. In Section 3, we recall the general variational inference setup. Then, we prove the main results in Section 4: we provide a reminder on the theory of distributions in Section 4.1, which is necessary to prove the main theorem in Section 4.2, and then, we provide some corollaries concerned with sums of penalties in Section 4.3. Afterwards, in Section 5, we show how our theoretical results can be used in a series of examples. We also provide in Section 6 an application of the preceding results to the construction of a heuristic for the strength of the penalty. Finally, we discuss the hypotheses and possible extensions of the main theoretical results in Section 7.

## 2 Related Work

Interpretations of existing empirical deep learning methods in a Bayesian variational inference framework include, for instance, *variational drop-out* (Kingma et al., 2015), a version of drop-out which fits the Bayesian framework. Further developments of variational drop-out have been made by Molchanov et al. (2017) for weight pruning and by Louizos et al. (2017) for neuron pruning.

Closer to our present work, the links between a penalized loss and the Bayesian point of view have previously been mentioned by Olshausen and Field (1997) and used to favor sparsity, but only within the MAP approximation. MacKay (1992b) noted the equivalence of the penalized loss  $\mathcal{L}(\boldsymbol{\theta}) = \ell(\boldsymbol{\theta}) + r(\boldsymbol{\theta})$  and the MAP loss  $\mathcal{L}_{\text{MAP}}(\boldsymbol{\theta}) = \ell(\boldsymbol{\theta}) - \ln(\alpha(\boldsymbol{\theta}))$  when  $\ell(\boldsymbol{\theta})$  is the negative log-likelihood of the training dataset given the parameters  $\boldsymbol{\theta}$ , and with a prior

$\alpha(\boldsymbol{\theta}) \propto \exp(-r(\boldsymbol{\theta}))$ . That is, finding the vector  $\hat{\boldsymbol{\theta}}$  minimizing a loss  $\mathcal{L}$  can be equivalent to finding the MAP estimator  $\hat{\boldsymbol{\theta}}_{\text{MAP}}$  by minimizing the loss  $\mathcal{L}_{\text{MAP}}$  with a well-chosen prior distribution  $\alpha$ . In the same vein, [Gribonval \(2011\)](#) has studied the Bayesian interpretation of the penalized least squares problem, and found out that a single penalty term can be interpreted as the influence of different priors, depending on the choice of the estimator.<sup>1</sup>

However, the MAP framework is not completely satisfying from a Bayesian point of view: instead of returning a distribution over the parameters, which contains information about their uncertainty, it returns the most reasonable value. In order to evaluate this uncertainty, [MacKay \(1992a\)](#) proposed a second-order approximation of the Bayesian posterior. In the process, [MacKay \(1995\)](#) also proposed a complete Bayesian framework and interpretation of neural networks. Still, this approximation of the Bayesian posterior is quite limited. In the same period, [Hinton and van Camp \(1993\)](#) applied the Minimum Description Length (MDL) principle to neural networks. Then, [MacKay \(2003\)](#) made the link between the MDL principle and variational inference, and [Graves \(2011\)](#) applied it to neural networks, allowing for variational approximations of the Bayesian posterior in a tractable way.

In a different perspective, variational inference is a key ingredient of the Variational Auto-Encoders (VAE) ([Kingma and Welling, 2013](#)) and subsequent variations and improvements, such as generative diffusion models ([Ho et al., 2020](#)). Given a distribution  $P$  of data points lying in a high-dimensional space (e.g., space of images), VAE models are trained to embed their inputs into a low-dimensional space (*encoding of the inputs*) and reconstruct these inputs by using only their embeddings (*decoding of the embeddings*). Such models can then be used to: 1) compress samples from  $P$  by encoding them; 2) generate new samples from  $P$  by decoding random noise. In the VAE setup, the loss is penalized by a KL-divergence between the distribution of the embeddings and a Bayesian prior, similarly to Eqn. (1). In this context, replacing the KL penalty term by a custom penalty is still under discussion ([Rombach et al., 2022](#), Appendix G). So, identifying conditions that a custom penalty must fulfill to correspond to the influence of a Bayesian prior remains a problem to be solved in the VAE training framework.

Moreover, variational inference is a way of approximating the Bayesian posterior which presents several advantages. First, as recalled by [Knoblauch et al. \(2022\)](#), the penalty used in variational inference is optimal for approximating the true Bayesian posterior. That is, any other penalty performs worse. Second, variational inference allows us to compute the variational posterior by gradient descent, which is a major advantage when the model to train is a neural network.

Additionally, a series of recent papers ([Alquier et al., 2020](#); [Chérif-Abdellatif, 2020](#); [Zhang and Gao, 2020](#)) provide theoretical results about the convergence speed of the variational posterior, according to the size of the training set. Therefore, translating user-specified penalties into Bayesian priors is a

---

<sup>1</sup>[Gribonval \(2011\)](#) focused on the MAP and the Minimum Mean-Square Error (MMSE) estimators.

way to fit a framework in which several results are already known.

### 3 Variational Inference

We propose here a reminder on variational inference, following Graves (2011).

From a Bayesian viewpoint, we describe the vector of parameters  $\boldsymbol{\theta} \in \mathbb{R}^N$  of a model as a random variable. Given a dataset  $\mathcal{D}$ , we denote by  $p_{\boldsymbol{\theta}}(\mathcal{D})$  the likelihood of  $\mathcal{D}$  given  $\boldsymbol{\theta}$ , that is, the probability that the model with parameter  $\boldsymbol{\theta}$  assigns to  $\mathcal{D}$ . For instance, with a dataset  $\mathcal{D} = \{(x_1, y_1), \dots, (x_n, y_n)\}$  of  $n$  input-output pairs and a model that assigns probabilities  $p_{\boldsymbol{\theta}}(y_i|x_i)$  for the outputs, then  $\ln p_{\boldsymbol{\theta}}(\mathcal{D}) = \sum_{i=1}^n \ln p_{\boldsymbol{\theta}}(y_i|x_i)$  is the total log-likelihood of the data given the model.

Given the dataset  $\mathcal{D}$ , the posterior distribution over the parameters  $\boldsymbol{\theta}$  is:

$$\pi_{\mathcal{D}}(\boldsymbol{\theta}) = \frac{p_{\boldsymbol{\theta}}(\mathcal{D})\alpha(\boldsymbol{\theta})}{\mathbb{P}(\mathcal{D})}, \quad \mathbb{P}(\mathcal{D}) = \int_{\mathbb{R}^N} p_{\boldsymbol{\theta}}(\mathcal{D})\alpha(\boldsymbol{\theta}) d\boldsymbol{\theta},$$

which is analytically intractable for complicated models, such as multi-layer nonlinear neural networks. However, the posterior  $\pi_{\mathcal{D}}$  can be approximated by looking for probability distributions  $\beta$  that minimize the loss:

$$\mathcal{L}_{\text{VI}}(\beta) = -\mathbb{E}_{\boldsymbol{\theta} \sim \beta} \ln p_{\boldsymbol{\theta}}(\mathcal{D}) + \text{KL}(\beta \parallel \alpha), \quad (2)$$

where  $\text{KL}(\beta \parallel \alpha) = \int_{\mathbb{R}^N} \ln\left(\frac{\beta(\boldsymbol{\theta})}{\alpha(\boldsymbol{\theta})}\right)\beta(\boldsymbol{\theta}) d\boldsymbol{\theta}$  is the Kullback–Leibler divergence. Indeed, one has  $\mathcal{L}_{\text{VI}}(\beta) = -\ln \mathbb{P}(\mathcal{D}) + \text{KL}(\beta \parallel \pi_{\mathcal{D}})$ , which is minimized with  $\beta = \pi_{\mathcal{D}}$ . Besides, one can decompose the KL-divergence into the *entropy* term and the *cross-entropy* term:

$$\text{KL}(\beta \parallel \alpha) = - \underbrace{\left[ - \int_{\mathbb{R}^N} \ln(\beta(\boldsymbol{\theta}))\beta(\boldsymbol{\theta}) d\boldsymbol{\theta} \right]}_{\text{entropy Ent}(\beta) \text{ of } \beta} + \underbrace{\left[ - \int_{\mathbb{R}^N} \ln(\alpha(\boldsymbol{\theta}))\beta(\boldsymbol{\theta}) d\boldsymbol{\theta} \right]}_{\text{cross-entropy between } \beta \text{ and } \alpha}.$$

The first term in the loss (2) represents the error made over the dataset  $\mathcal{D}$ : it is small if  $\beta$  is concentrated around good parameters  $\boldsymbol{\theta}$ . The second term can be seen as a user-defined penalty over  $\beta$  that keeps it from diverging too much from the prior  $\alpha$ . Moreover, for any distribution  $\beta$ , the quantity  $\mathcal{L}_{\text{VI}}(\beta)$  is an upper bound on the Bayesian negative log-likelihood of the data:  $\mathcal{L}_{\text{VI}}(\beta) \geq -\ln \int_{\mathbb{R}^N} \alpha(\boldsymbol{\theta})p_{\boldsymbol{\theta}}(\mathcal{D}) d\boldsymbol{\theta}$  (Jordan et al., 1999).

In variational inference, a parametric family  $(\beta_{\mathbf{u}})_{\mathbf{u}}$  of probability distributions is fixed, and one looks for the best approximation  $\beta_{\mathbf{u}^*}$  of the Bayesian posterior by minimizing  $\mathcal{L}_{\text{VI}}(\beta_{\mathbf{u}})$  according to  $\mathbf{u}$ . So,  $(\beta_{\mathbf{u}})_{\mathbf{u}}$  is the *variational family*,  $\mathbf{u}$  is the *variational parameter* and  $\beta_{\mathbf{u}^*}$  is the *variational posterior*. Importantly, for some families such as Gaussians with fixed variance, the gradient of  $\mathcal{L}_{\text{VI}}(\beta_{\mathbf{u}})$  can be computed if the gradients of  $\ln p_{\boldsymbol{\theta}}(\mathcal{D})$  can be computed (Graves, 2011), so that  $\mathcal{L}_{\text{VI}}(\beta_{\mathbf{u}})$  can be optimized by stochastic gradient descent. Thus, this is well-suited for models such as neural networks.

To summarize, we consider a parametric family  $(\beta_{\mathbf{u}})_{\mathbf{u}}$  with parameter  $\mathbf{u}$ , where each  $\beta_{\mathbf{u}}$  is a probability distribution over  $\boldsymbol{\theta} \in \mathbb{R}^N$ . Then we learn the variational parameters  $\mathbf{u}$  instead of  $\boldsymbol{\theta}$ . For instance, we can choose one of the following families of variational posteriors.

**Example 1.** *The family of products of Gaussian distributions over  $\boldsymbol{\theta} = (\theta_1, \dots, \theta_N)$ :*

$$\beta_{\mathbf{u}} = \beta_{\boldsymbol{\mu}, \boldsymbol{\sigma}} = \mathcal{N}(\mu_1, \sigma_1^2) \otimes \dots \otimes \mathcal{N}(\mu_N, \sigma_N^2).$$

*In this case, the parameters of the model estimated by VI are random and independently sampled from different Gaussian distributions  $\mathcal{N}(\mu_k, \sigma_k^2)$ . Instead of learning them directly, the vector of parameters  $\mathbf{u} = (\mu_1, \sigma_1, \dots, \mu_N, \sigma_N)$  is learned in such a way that they minimize  $\mathcal{L}_{\text{VI}}(\beta_{\mathbf{u}})$ .*

**Example 2.** *The family of products of Dirac distributions over  $\boldsymbol{\theta} = (\theta_1, \dots, \theta_N)$ :*

$$\beta_{\mathbf{u}} = \beta_{\boldsymbol{\mu}} = \delta_{\mu_1} \otimes \dots \otimes \delta_{\mu_N}.$$

*In this case, the parameters are deterministic:  $\theta_k$  and  $\mu_k$  are identical for all  $k$ .*

## 4 Bayesian Interpretation of Penalties

In this section, we provide conditions ensuring that a penalty  $r(\mathbf{u})$  over the parameters of a variational posterior can be interpreted as a Kullback–Leibler divergence with respect to a prior  $\alpha$ ; namely, that:

$$\exists K \in \mathbb{R}, \forall \mathbf{u}, \quad r(\mathbf{u}) + K = \text{KL}(\beta_{\mathbf{u}} \parallel \alpha), \quad (3)$$

where the constant  $K$  is irrelevant for optimization. In the process, we give a formula expressing  $\alpha$  as a function of  $r(\cdot)$ .

More precisely, under some conditions over the variational family (Conditions (a)-(c)) and the search space of the prior (space  $\mathcal{T}(\mathbb{R}^N)$ ), we provide a set of necessary and sufficient conditions (Conditions (i)-(iv)) that the penalty must fulfill to ensure existence and uniqueness of a corresponding prior. The proof is given in Section 4.2 and is divided into two parts: in Theorem 1, we state some regularity conditions on the variational family  $(\beta_{\mathbf{u}})_{\mathbf{u}}$  and the penalty  $r(\mathbf{u})$ , so that a candidate prior  $\exp(A_{\mathbf{u}})$  would emerge; in Corollary 1, we state conditions over  $A_{\mathbf{u}}$  so that  $\exp(A_{\mathbf{u}})$  corresponds effectively to the density of a prior. Then, in Section 4.3, we prove that the prior exhibits a specific structure when the corresponding penalty is made of a sum of penalties, each one over a different subset of the parameters.

Since some notions of theory of distributions are necessary to our derivation, we start by providing a reminder.

## 4.1 Reminder on Theory of Distributions

In order to explain the main result, we recall some basic concepts of theory of distributions. We use three functional spaces: the Schwartz class  $\mathcal{S}(\mathbb{R}^N)$ , the space of tempered distributions  $\mathcal{S}'(\mathbb{R}^N)$ , and the space of distributions with compact support  $\mathcal{E}'(\mathbb{R}^N)$ .

We recall the definition of the space of distributions  $\mathcal{D}'(\mathbb{R}^N)$ . We denote by  $C^\infty(\mathbb{R}^N)$  the space of infinitely derivable functions mapping  $\mathbb{R}^N$  to  $\mathbb{R}$ , and by  $C_c^\infty(\mathbb{R}^N) \subset C^\infty(\mathbb{R}^N)$  the subspace of functions with compact support, that is  $\varphi \in C_c^\infty(\mathbb{R}^N)$  if, and only if:

$$\varphi \in C^\infty(\mathbb{R}^N) \text{ and } \exists K \subset \mathbb{R}^N \text{ compact s.t.: } \{x \in \mathbb{R}^N : \varphi(x) \neq 0\} \subseteq K.$$

The set  $\{x \in \mathbb{R}^N : \varphi(x) \neq 0\}$  is also denoted by  $\text{supp}(\varphi)$ .

**Space of distributions  $\mathcal{D}'(\mathbb{R}^N)$ .** The space of distributions  $\mathcal{D}'(\mathbb{R}^N)$  is defined as the space of continuous linear forms over  $C_c^\infty(\mathbb{R}^N)$ . For any *distribution*  $T \in \mathcal{D}'(\mathbb{R}^N)$ , we denote by  $\langle T, \phi \rangle$  the value of  $T$  at a given *test function*  $\varphi \in C_c^\infty(\mathbb{R}^N)$ .

More formally,  $T \in \mathcal{D}'(\mathbb{R}^N)$  if, and only if, for all compact set  $K$  of  $\mathbb{R}^N$ , there exists  $p \in \mathbb{N}$  and  $C > 0$  such that:

$$\begin{aligned} \forall \varphi \in C_c^\infty(\mathbb{R}^N) \text{ with } \text{supp}(\varphi) \subseteq K, \\ |\langle T, \varphi \rangle| \leq C \sup_{|\alpha| \leq p} \|\partial^\alpha \varphi\|_\infty, \end{aligned}$$

where  $\alpha \in \mathbb{N}^N$  is a multi-index representing the order of derivation according to each of the  $N$  variables, and  $|\alpha| = \sum_{i=1}^N \alpha_i$ .

Distributions are easier to visualize in specific cases. For instance, if a function  $f : \mathbb{R}^N \rightarrow \mathbb{R}$  is integrable on every compact set  $K \in \mathbb{R}^N$ , then we can define a distribution  $T_f$  by:

$$\forall \varphi \in C_c^\infty(\mathbb{R}^N), \quad \langle T_f, \varphi \rangle = \int f \varphi.$$

Therefore, distributions are often called “generalized functions”, since most functions can be seen as distributions. In fact, by abuse of notation,  $\langle f, \varphi \rangle$  stands for  $\langle T_f, \varphi \rangle$ .

Another classic example of distribution is the Dirac at zero  $\delta$ , defined as follows:

$$\forall \varphi \in C_c^\infty(\mathbb{R}^N), \quad \langle \delta, \varphi \rangle = \varphi(0).$$

Finally, the derivatives  $\partial^\alpha T$  of a distribution  $T$  are defined as follows:

$$\forall \varphi \in C_c^\infty(\mathbb{R}^N), \quad \langle \partial^\alpha T, \varphi \rangle = (-1)^{|\alpha|} \langle T, \partial^\alpha \varphi \rangle.$$

Notably, when  $N = 1$ ,  $\langle T', \varphi \rangle = -\langle T, \varphi' \rangle$ .

**Schwartz class  $\mathcal{S}(\mathbb{R}^N)$ .**  $\varphi$  belongs to the Schwartz class  $\mathcal{S}(\mathbb{R}^N)$ , if and only if,  $\varphi \in C^\infty(\mathbb{R}^N)$  and:

$$\forall p \in \mathbb{N}, \exists C_p > 0 : \sup_{|\alpha| \leq p, |\beta| \leq p} \|x^\alpha \partial^\beta \varphi(x)\|_\infty \leq C_p.$$

In other words,  $\mathcal{S}(\mathbb{R}^N)$  contains smooth and rapidly decreasing functions. For instance, any Gaussian density function belongs to  $\mathcal{S}(\mathbb{R}^N)$ .

We can easily define the Fourier transform on  $\mathcal{S}(\mathbb{R}^N)$ . Let  $\varphi \in \mathcal{S}(\mathbb{R}^N)$ :

$$(\mathcal{F}\varphi)(\xi) = \int_{\mathbb{R}^N} \varphi(x) e^{-i\xi x} dx.$$

Thus,  $\mathcal{F}^{-1} = (2\pi)^{-N} \bar{\mathcal{F}}$ , where  $\bar{\mathcal{F}}\varphi = \mathcal{F}\check{\varphi}$  and  $\check{\varphi}(x) = \varphi(-x)$ .

**Space of tempered distributions  $\mathcal{S}'(\mathbb{R}^N)$ .** Let  $T \in \mathcal{D}'(\mathbb{R}^N)$  be a distribution.  $T$  belongs to  $\mathcal{S}'(\mathbb{R}^N)$  if, and only if, there exists  $p \in \mathbb{N}$  and  $C > 0$  such that:

$$\forall \varphi \in C_c^\infty(\mathbb{R}^N), |\langle T, \varphi \rangle| \leq \sup_{|\alpha| \leq p, |\beta| \leq p} \|x^\alpha \partial^\beta \varphi\|_\infty.$$

The Fourier transform is defined on  $\mathcal{S}'(\mathbb{R}^N)$  by duality. For any  $T \in \mathcal{S}'(\mathbb{R}^N)$ ,  $\mathcal{F}T \in \mathcal{S}'(\mathbb{R}^N)$  and is defined by:

$$\forall \varphi \in \mathcal{S}(\mathbb{R}^N), \langle \mathcal{F}T, \varphi \rangle = \langle T, \mathcal{F}\varphi \rangle.$$

Notably, this definition allows us to compute the Fourier transform of functions that do not lie in  $\mathcal{L}^2$ . This is very useful in the applications of Theorem 1, where we need the Fourier transform of  $f : x \mapsto x^2$ , which is  $\mathcal{F}f = -2\pi\delta'' \in \mathcal{S}'(\mathbb{R})$  (where  $\delta''$  is the second derivative of the Dirac, defined below).

**Space of distributions with compact support  $\mathcal{E}'(\mathbb{R}^N)$ .** Let  $T \in \mathcal{D}'(\mathbb{R}^N)$ . The support of  $T$  is defined by:

$$\text{supp}(T) = \mathbb{R}^N \setminus \{x \in \mathbb{R}^N : \exists \omega \in \mathcal{N}_x \text{ s.t. } T|_\omega = 0\},$$

where  $\mathcal{N}_x$  is the set of neighborhoods of  $x$ . Thus,  $T$  is said to have a compact support if, and only if,  $\text{supp}(T)$  is contained into a compact subset of  $\mathbb{R}^N$ . As fundamental property of  $\mathcal{E}'(\mathbb{R}^N)$ , one should notice that  $\mathcal{E}'(\mathbb{R}^N) \subset \mathcal{S}'(\mathbb{R}^N)$ . That is, the Fourier transform is defined on  $\mathcal{E}'(\mathbb{R}^N)$ .

For instance, the Dirac at zero  $\delta$  and its derivatives  $\delta^{(k)}$  have support  $\{0\}$ , which is compact. For any test function  $\varphi \in C_c^\infty(\mathbb{R})$ :

$$\langle \delta, \varphi \rangle = \varphi(0), \quad \langle \delta^{(k)}, \varphi \rangle = (-1)^k \varphi^{(k)}(0).$$

## 4.2 Main Results

**Notation and statement.** In order to approximate the posterior distribution of a vector  $\boldsymbol{\theta} \in \mathbb{R}^N$ , we denote by  $(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})_{\boldsymbol{\mu}, \boldsymbol{\nu}}$  the family of variational posteriors over  $\boldsymbol{\theta}$ , parameterized by their mean  $\boldsymbol{\mu}$  and a vector of additional parameters  $\boldsymbol{\nu}$ . The basic example is a multivariate Gaussian distribution  $\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}}$  parameterized by its mean  $\boldsymbol{\mu} \in \mathbb{R}^N$  and its covariance matrix  $\boldsymbol{\nu} = \boldsymbol{\Sigma} \in \mathbb{R}^{N \times N}$ .

We say that the family  $(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})_{\boldsymbol{\mu}, \boldsymbol{\nu}}$  of variational posteriors is *translation-invariant* if  $\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}}(\boldsymbol{\theta}) = \beta_{0, \boldsymbol{\nu}}(\boldsymbol{\theta} - \boldsymbol{\mu})$  for all  $\boldsymbol{\mu}, \boldsymbol{\nu}, \boldsymbol{\theta}$ .

We denote by  $\text{KL}(f\|g)$  the KL-divergence extended to functions  $f$  and  $g$ , which are not necessarily probability densities:

$$\text{KL}(f\|g) = \int \ln \left( \frac{f(\boldsymbol{\theta})}{g(\boldsymbol{\theta})} \right) f(\boldsymbol{\theta}) \, d\boldsymbol{\theta}$$

We denote by  $r_{\boldsymbol{\nu}}(\boldsymbol{\mu})$  some penalty, applied to the distribution  $\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}}$ .

In the following, we always restrict ourselves to priors  $\alpha \in \mathcal{T}(\mathbb{R}^N) = \{\alpha \text{ s.t. } \ln(\alpha) \in \mathcal{S}'(\mathbb{R}^N)\}$ , i.e. log-tempered probability distributions, hence the condition  $\ln(\alpha) \in \mathcal{S}'(\mathbb{R}^N)$  in the results below.

Also, we make the following assumptions on the variational family:

**Assumption 1.** Let  $(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})_{\boldsymbol{\mu}, \boldsymbol{\nu}}$  be a family of variational posteriors such that:

- (a)  $(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})_{\boldsymbol{\mu}, \boldsymbol{\nu}}$  is translation-invariant;
- (b) for all  $\boldsymbol{\nu}$ , the density  $\beta_{0, \boldsymbol{\nu}}$  has finite entropy and belongs to  $\mathcal{S}(\mathbb{R}^N)$ ;
- (c) for all  $\boldsymbol{\nu}$ ,  $\mathcal{F}\beta_{0, \boldsymbol{\nu}}$  is nonzero everywhere.

**Theorem 1.** We assume that the variational family  $(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})_{\boldsymbol{\mu}, \boldsymbol{\nu}}$  fulfills Assumption 1. Let  $r_{\boldsymbol{\nu}}(\boldsymbol{\mu})$  be a penalty over  $\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}}$ . We assume that, for all  $\boldsymbol{\nu}$ :

- (i)  $r_{\boldsymbol{\nu}} \in \mathcal{S}'(\mathbb{R}^N)$ ;
- (ii) the quotient  $\mathcal{F}r_{\boldsymbol{\nu}}/\mathcal{F}\check{\beta}_{0, \boldsymbol{\nu}}$  belongs to  $\mathcal{S}'(\mathbb{R}^N)$ , where  $\check{\beta}_{0, \boldsymbol{\nu}}(\boldsymbol{\theta}) := \beta_{0, \boldsymbol{\nu}}(-\boldsymbol{\theta})$ .

Then, there exists a unique function  $A_{\boldsymbol{\nu}} \in \mathcal{S}'(\mathbb{R}^N)$  such that:

$$r_{\boldsymbol{\nu}}(\boldsymbol{\mu}) = \text{KL}(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}} \parallel \exp(A_{\boldsymbol{\nu}})), \quad (4)$$

$$\text{moreover:} \quad A_{\boldsymbol{\nu}} = -\text{Ent}(\beta_{0, \boldsymbol{\nu}}) \mathbb{1} - \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_{\boldsymbol{\nu}}}{\mathcal{F}\check{\beta}_{0, \boldsymbol{\nu}}} \right], \quad (5)$$

where  $\check{\beta}_{0, \boldsymbol{\nu}}(\boldsymbol{\theta}) := \beta_{0, \boldsymbol{\nu}}(-\boldsymbol{\theta})$  and  $\mathbb{1}$  is the constant function equal to 1.

*Proof.* The proof is based on the resolution of a classical integral equation (Polyanin and Manzhirov, 1998, Section 10.3-1) adapted to the wider framework of the theory of distributions. This extension is necessary, since the Fourier transform of the widely-used  $\mathcal{L}^2$  penalty  $r(x) = x^2$  is  $-2\pi\delta''$ , which is a distribution that cannot be expressed as a function.

Let  $A_\nu$  be a function in  $\mathcal{S}'(\mathbb{R}^N)$ . We have:

$$\begin{aligned} \text{KL}(\beta_{\mu,\nu} \parallel \exp(A_\nu)) &= -\text{Ent}(\beta_{0,\nu}) - \int_{\mathbb{R}^N} A_\nu(\boldsymbol{\theta}) \beta_{\mu,\nu}(\boldsymbol{\theta}) \, d\boldsymbol{\theta} \\ &= -\text{Ent}(\beta_{0,\nu}) - \int_{\mathbb{R}^N} A_\nu(\boldsymbol{\theta}) \check{\beta}_{0,\nu}(\boldsymbol{\mu} - \boldsymbol{\theta}) \, d\boldsymbol{\theta} \\ &= -\text{Ent}(\beta_{0,\nu}) - (A_\nu * \check{\beta}_{0,\nu})(\boldsymbol{\mu}), \end{aligned}$$

which exists since the convolution between  $\check{\beta}_{0,\nu} \in \mathcal{S}(\mathbb{R}^N)$  and  $A_\nu \in \mathcal{S}'(\mathbb{R}^N)$  is well-defined (Stein and Weiss, 2016, Chap. I, Theorem 3.13).

Since  $\mathcal{F}$  is an isomorphism on  $\mathcal{S}'(\mathbb{R}^N)$  (Stein and Weiss, 2016, Chap. I.3), we have:

$$r_\nu(\cdot) = \text{KL}(\beta_{\cdot,\nu} \parallel \exp(A_\nu)) \quad \Leftrightarrow \quad \mathcal{F}r_\nu = \mathcal{F}[\text{KL}(\beta_{\cdot,\nu} \parallel \exp(A_\nu))].$$

Besides, we have:

$$\mathcal{F}[\text{KL}(\beta_{\cdot,\nu} \parallel \exp(A_\nu))] = -(2\pi)^N \text{Ent}(\beta_{0,\nu})\delta - (\mathcal{F}A_\nu) \cdot (\mathcal{F}\check{\beta}_{0,\nu}),$$

by using  $\mathcal{F}\mathbb{1} = (2\pi)^N\delta$  and  $\mathcal{F}(A_\nu * \check{\beta}_{0,\nu}) = (\mathcal{F}A_\nu) \cdot (\mathcal{F}\check{\beta}_{0,\nu})$ , which holds since  $\check{\beta}_{0,\nu} \in \mathcal{S}(\mathbb{R}^N)$  and  $A_\nu \in \mathcal{S}'(\mathbb{R}^N)$ . See the proof of (Stein and Weiss, 2016, Chap. I, Theorem 3.18) for more details. So:

$$\begin{aligned} r_\nu(\cdot) = \text{KL}(\beta_{\cdot,\nu} \parallel \exp(A_\nu)) &\quad \Leftrightarrow \quad \mathcal{F}r_\nu = -(2\pi)^N \text{Ent}(\beta_{0,\nu})\delta - (\mathcal{F}A_\nu) \cdot (\mathcal{F}\check{\beta}_{0,\nu}) \\ &\quad \Leftrightarrow \quad \mathcal{F}A_\nu = -(2\pi)^N \text{Ent}(\beta_{0,\nu}) \frac{\delta}{\mathcal{F}\check{\beta}_{0,\nu}} - \frac{\mathcal{F}r_\nu}{\mathcal{F}\check{\beta}_{0,\nu}} \\ &\quad \Leftrightarrow \quad A_\nu = -\text{Ent}(\beta_{0,\nu})\mathbb{1} - \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_\nu}{\mathcal{F}\check{\beta}_{0,\nu}} \right]. \end{aligned} \tag{6}$$

$$\tag{7}$$

The equivalence with (6) holds since  $\mathcal{F}\check{\beta}_{0,\nu}$  is assumed to be nonzero everywhere (Condition (c)). We obtain (7) by using  $\frac{\delta}{\mathcal{F}\check{\beta}_{0,\nu}} = \frac{\delta}{(\mathcal{F}\check{\beta}_{0,\nu})(0)} = \delta$ , since  $\check{\beta}_{0,\nu}$  is a probability distribution, and  $\mathcal{F}^{-1}[(2\pi)^N\delta] = \mathbb{1}$ . To ensure the equivalence with (6), we need to check that  $\frac{\mathcal{F}r_\nu}{\mathcal{F}\check{\beta}_{0,\nu}} \in \mathcal{S}'(\mathbb{R}^N)$ , which is true by using Condition (ii), and we use the fact that  $\mathcal{F}^{-1}$  is an isomorphism on  $\mathcal{S}'(\mathbb{R}^N)$ .  $\square$

**Remarks about Theorem 1.** Theorem 1 holds under some technical assumptions. Even if some of them fail, the formula is still useful to compute a candidate prior  $\alpha$  from a penalty  $r(\cdot)$ . One would have to apply Equation (5) on a penalty  $r$  to compute  $A_\nu$ , then check that Conditions (iii) and (iv) hold, define the distribution  $\alpha \propto \exp(A)$  as in Eqn. (9), and finally compute  $\text{KL}(\beta_{\mu,\nu} \parallel \alpha)$  analytically and compare it to  $r$ . However, this process does not guarantee the uniqueness of the solution  $\alpha$ .

**Remark 1.** In Theorem 1, Condition (ii) can be replaced by Condition (ii'):

(ii')  $\mathcal{F}r$  belongs to  $\mathcal{E}'(\mathbb{R}^N)$ ,

where  $\mathcal{E}'(\mathbb{R}^N)$  is the space of distributions with compact support. Indeed:

Conditions (ii') and (c)  $\Rightarrow$  Condition (ii).

It is more convenient to use Condition (ii') when the penalty  $r$  is known to have a Fourier transform with compact support. For instance, for any integer  $p > 0$ , the Fourier transform of the penalty  $r_{2p}(x) = x^{2p}$  is:

$$\mathcal{F}r_{2p} = (-1)^p 2\pi \delta^{(2p)} \in \mathcal{E}'(\mathbb{R}),$$

where  $\delta^{(2p)}$  is the  $2p$ -derivative of the Dirac distribution, which has a compact support. Moreover, by linearity of  $\mathcal{F}$ , any *finite* linear combination of  $(r_{2p})_p$  has also a Fourier transform that is in  $\mathcal{E}'(\mathbb{R})$ .

**Consequences for the prior.** At this point, the situation is:

$$\begin{aligned} \text{problem to solve:} \quad & r_{\nu}(\boldsymbol{\mu}) + K = \text{KL}(\beta_{\boldsymbol{\mu}, \nu} \| \alpha), \\ \text{problem solved:} \quad & r_{\nu}(\boldsymbol{\mu}) = \text{KL}(\beta_{\boldsymbol{\mu}, \nu} \| \exp(A_{\nu})), \end{aligned}$$

where  $A_{\nu}$  is defined in Eqn. (5),  $K \in \mathbb{R}$  and  $\alpha$  is a probability distribution to find. Informally, we expect  $\alpha$  to be proportional to  $\exp(A_{\nu})$ . So, two conditions have to be taken into account:  $A_{\nu}$  should not depend on  $\nu$  and  $\exp(A_{\nu})$  should integrate to 1. In order to obtain such an  $\alpha$ , the two following results, although immediate, will be helpful.

**Lemma 1.** Let  $r_{\nu}^{(1)}(\boldsymbol{\mu})$  and  $r_{\nu}^{(2)}(\boldsymbol{\mu})$  be two penalties fulfilling the conditions of Theorem 1. We can build the solutions of Eqn. (4) for  $r_{\nu}^{(1)}$  and  $r_{\nu}^{(2)}$  by introducing  $\bar{A}_{\nu}^{(1)}$  and  $\bar{A}_{\nu}^{(2)}$ :

$$\begin{aligned} r_{\nu}^{(1)}(\boldsymbol{\mu}) &= \text{KL}(\beta_{\boldsymbol{\mu}, \nu} \| \exp[-\text{Ent}(\beta_{0, \nu}) - \bar{A}_{\nu}^{(1)}]), & \bar{A}_{\nu}^{(1)} &:= \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_{\nu}^{(1)}}{\mathcal{F}\check{\beta}_{0, \nu}} \right], \\ r_{\nu}^{(2)}(\boldsymbol{\mu}) &= \text{KL}(\beta_{\boldsymbol{\mu}, \nu} \| \exp[-\text{Ent}(\beta_{0, \nu}) - \bar{A}_{\nu}^{(2)}]), & \bar{A}_{\nu}^{(2)} &:= \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_{\nu}^{(2)}}{\mathcal{F}\check{\beta}_{0, \nu}} \right], \end{aligned}$$

Thus, for any  $K \in \mathbb{R}$ , we have:

$$r_{\nu}^{(1)}(\boldsymbol{\mu}) + Kr_{\nu}^{(2)}(\boldsymbol{\mu}) = \text{KL}(\beta_{\boldsymbol{\mu}, \nu} \| \exp[-\text{Ent}(\beta_{0, \nu}) - \bar{A}_{\nu}^{(1)} - K\bar{A}_{\nu}^{(2)}]).$$

**Lemma 2.** Specifically, Lemma 1 allows us to show that:

(1) for  $K = 1$  and  $r_{\nu}^{(2)}(\boldsymbol{\mu}) = f(\boldsymbol{\nu})$ , we have:

$$\begin{aligned} r_{\nu}^{(1)}(\boldsymbol{\mu}) + f(\boldsymbol{\nu}) &= \text{KL}(\beta_{\boldsymbol{\mu}, \nu} \| \exp[-\text{Ent}(\beta_{0, \nu}) - \bar{A}_{\nu}^{(1)} - f(\boldsymbol{\nu})]) \\ A_{\nu}(\boldsymbol{\theta}) &:= -\text{Ent}(\beta_{0, \nu}) - \bar{A}_{\nu}^{(1)}(\boldsymbol{\theta}) - f(\boldsymbol{\nu}). \end{aligned}$$

(2) for  $r_{\nu}^{(2)}(\boldsymbol{\mu}) = 1$ , we have:

$$\begin{aligned} r_{\nu}^{(1)}(\boldsymbol{\mu}) + K &= \text{KL}(\beta_{\boldsymbol{\mu}, \nu} \| \exp[-\text{Ent}(\beta_{0, \nu}) - \bar{A}_{\nu}^{(1)} - K]) \\ A_{\nu}(\boldsymbol{\theta}) &:= -\text{Ent}(\beta_{0, \nu}) - \bar{A}_{\nu}^{(1)}(\boldsymbol{\theta}) - K. \end{aligned}$$

In Lemma 2, Result (1) can be used to build a function  $f(\nu)$  in order to make the  $A_{\nu}$  independent from  $\nu$ , which is necessary to consider  $A_{\nu}$  as a log-prior. This change comes at a cost: the initial penalty  $r_{\nu}^{(1)}(\boldsymbol{\mu})$  should be balanced by a non-constant term  $f(\nu)$ . Result (2) can be used to tune the constant  $K$  in such a way that  $\exp(A_{\nu})$  integrates to 1. As a result,  $K$  is added to the penalty  $r$ , which does not influence the optimization process. But  $\exp(-\text{Ent}(\beta_{0, \nu}) - \bar{A}_{\nu}^{(1)})$  must be integrable.

Also, it is worth noticing that, if Eqn. (4) admits a solution for  $r_{\nu}$ , then it admits a solution for  $r_{\nu} + K$ , with any  $K \in \mathbb{R}$ .

**Corollary 1.** *We assume that the variational family  $(\beta_{\boldsymbol{\mu}, \nu})_{\boldsymbol{\mu}, \nu}$  fulfills Assumption 1. We define the following conditions on the penalty and the variational family:*

- (i)  $r_{\nu} \in \mathcal{S}'(\mathbb{R}^N)$ ;
- (ii) the quotient  $\mathcal{F}r_{\nu}/\mathcal{F}\check{\beta}_{0, \nu}$  belongs to  $\mathcal{S}'(\mathbb{R}^N)$ , where  $\check{\beta}_{0, \nu}(\boldsymbol{\theta}) := \beta_{0, \nu}(-\boldsymbol{\theta})$ ;
- (iii)  $A_{\nu} := -\text{Ent}(\beta_{0, \nu}) \mathbb{1} - \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_{\nu}}{\mathcal{F}\check{\beta}_{0, \nu}} \right]$  does not depend on  $\nu$ ; let  $A := A_{\nu}$ ;
- (iv)  $A$  is such that  $\int \exp(A) < \infty$ .

If Conditions (i)-(iv) are fulfilled, then there exists a unique probability distribution  $\alpha \in \mathcal{T}(\mathbb{R}^N)$  such that:

$$\exists K \in \mathbb{R} : \quad r_{\nu}(\boldsymbol{\mu}) + K = \text{KL}(\beta_{\boldsymbol{\mu}, \nu} \| \alpha), \quad (8)$$

$$\text{moreover:} \quad \alpha(\boldsymbol{\theta}) = \frac{\exp(A(\boldsymbol{\theta}))}{\int \exp(A(\mathbf{t})) \, d\mathbf{t}}. \quad (9)$$

Incidentally,  $K = \ln(\int \exp(A(\mathbf{t})) \, d\mathbf{t})$ .

Conversely, if  $\alpha \in \mathcal{T}(\mathbb{R}^N)$  is a probability distribution, and  $K$  a constant, then the penalty  $r_{\nu}(\boldsymbol{\mu}) := \text{KL}(\beta_{\boldsymbol{\mu}, \nu} \| \alpha) - K$  fulfills Conditions (i)-(iv).

### 4.3 Sum of Penalties

In the preceding section, we have established a result in the general case of a penalty over the entire vector of means  $\boldsymbol{\mu}$ . In practice, it is common (Shalev-Shwartz and Ben-David, 2014, Chap. 13) to deal with a penalty that can be decomposed into a sum of sub-penalties, each one over a subset of the means  $(\mu_1, \dots, \mu_N)$ . Also, this framework encompasses the ‘‘weight decay’’ technique (Krogh and Hertz, 1992), widely used in deep learning, which consists in using a penalty proportional to the sum of the squared parameters of the model. In

this section, we show how such structure of the penalty affects the structure of the corresponding prior. In the simplest case when a penalty can be written as a sum of  $N$  penalties, each one over a component  $\mu_i$  of  $\boldsymbol{\mu}$ , then the corresponding prior  $\alpha$  over  $\boldsymbol{\theta} \in \mathbb{R}^N$  is a product of  $N$  distributions:  $\alpha(\boldsymbol{\theta}) = \alpha_1(\theta_1) \cdots \alpha_N(\theta_N)$ .

**Proposition 1.** *Let  $(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})_{\boldsymbol{\mu}, \boldsymbol{\nu}}$  a variational family fulfilling Assumption 1, with parameters  $\boldsymbol{\mu} \in \mathbb{R}^N$  and  $\boldsymbol{\nu}$ . Let  $(I_i)_{1 \leq i \leq n}$  be a partition of  $\{1, \dots, N\}$ . Given a vector  $\boldsymbol{\theta} \in \mathbb{R}^N$ , we denote by  $\boldsymbol{\theta}_{I_i}$  the vector of  $\mathbb{R}^{|I_i|}$  containing only the components of  $\boldsymbol{\theta}$  whose index belongs to  $I_i$ .*

We consider the following penalty on  $\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}}$ :

$$r_{\boldsymbol{\nu}}(\boldsymbol{\mu}) := \sum_{i=1}^n r_{\boldsymbol{\nu}}^{(i)}(\boldsymbol{\mu}_{I_i}).$$

Thus, under Conditions (i) and (ii), the solution of Eqn. (4) is:

$$A_{\boldsymbol{\nu}}(\boldsymbol{\theta}) = -\text{Ent}(\beta_{0, \boldsymbol{\nu}}) - \sum_{k=1}^n \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_{\boldsymbol{\nu}}^{(i)}}{\mathcal{F}\check{\beta}_{I_i|0, \boldsymbol{\nu}}} \right] (\boldsymbol{\theta}_{I_i}), \quad (10)$$

where  $\beta_{I_i|0, \boldsymbol{\nu}}(\boldsymbol{\theta}_{I_i}) = \int \beta_{0, \boldsymbol{\nu}}(\boldsymbol{\theta}_{I_1}, \dots, \boldsymbol{\theta}_{I_n}) d\boldsymbol{\theta}_{I_1} \cdots d\boldsymbol{\theta}_{I_{i-1}} d\boldsymbol{\theta}_{I_{i+1}} \cdots d\boldsymbol{\theta}_{I_n}$  is the marginal probability of  $\beta_{0, \boldsymbol{\nu}}$  at indices  $I_i$ .

*Proof.* According to Theorem 1, the solution of Eqn. (4) is:

$$A_{\boldsymbol{\nu}} = -\text{Ent}(\beta_{0, \boldsymbol{\nu}}) \mathbb{1} - \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_{\boldsymbol{\nu}}}{\mathcal{F}\check{\beta}_{0, \boldsymbol{\nu}}} \right].$$

We have  $r_{\boldsymbol{\nu}}(\boldsymbol{\theta}) = \sum_{i=1}^n r_{\boldsymbol{\nu}}^{(i)}(\boldsymbol{\theta}_{I_i}) \mathbb{1}(\boldsymbol{\theta})$ , where  $\mathbb{1}$  is the constant function equal to 1 on  $\mathbb{R}^N$ . So,  $\mathcal{F}r_{\boldsymbol{\nu}} = \sum_{i=1}^n (2\pi)^{N-|I_i|} \mathcal{F}r_{\boldsymbol{\nu}}^{(i)} \delta|_{\bar{I}_i}$ , where  $\delta|_{\bar{I}_i}$  is the Dirac distribution on functions of  $(\boldsymbol{\theta}_{I_1}, \dots, \boldsymbol{\theta}_{I_{i-1}}, \boldsymbol{\theta}_{I_{i+1}}, \dots, \boldsymbol{\theta}_{I_n})$ .

Thus, we have:

$$\frac{\mathcal{F}r_{\boldsymbol{\nu}}}{\mathcal{F}\check{\beta}_{0, \boldsymbol{\nu}}} = \sum_{i=1}^n (2\pi)^{N-|I_i|} \frac{\mathcal{F}r_{\boldsymbol{\nu}}^{(i)} \delta|_{\bar{I}_i}}{\mathcal{F}\check{\beta}_{0, \boldsymbol{\nu}}}.$$

Let us compute the right-hand side. Let  $\varphi \in C_c^\infty(\mathbb{R}^N)$  be a test function. We have:

$$\begin{aligned} \left\langle \frac{\mathcal{F}r_{\boldsymbol{\nu}}^{(i)} \delta|_{\bar{I}_i}}{\mathcal{F}\check{\beta}_{0, \boldsymbol{\nu}}}, \varphi \right\rangle &= \left\langle \mathcal{F}r_{\boldsymbol{\nu}}^{(i)}, \frac{\varphi(0, \dots, 0, \cdot, 0, \dots, 0)}{(\mathcal{F}\check{\beta}_{0, \boldsymbol{\nu}})(0, \dots, 0, \cdot, 0, \dots, 0)} \right\rangle \\ &= \left\langle \mathcal{F}r_{\boldsymbol{\nu}}^{(i)}, \frac{\varphi(0, \dots, 0, \cdot, 0, \dots, 0)}{\mathcal{F}\check{\beta}_{I_i|0, \boldsymbol{\nu}}} \right\rangle = \left\langle \frac{\mathcal{F}r_{\boldsymbol{\nu}}^{(i)}}{\mathcal{F}\check{\beta}_{I_i|0, \boldsymbol{\nu}}} \delta|_{\bar{I}_i}, \varphi \right\rangle. \end{aligned}$$

Therefore:

$$\mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_{\boldsymbol{\nu}}}{\mathcal{F}\check{\beta}_{0, \boldsymbol{\nu}}} \right] (\boldsymbol{\theta}) = \sum_{i=1}^n \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_{\boldsymbol{\nu}}^{(i)}}{\mathcal{F}\check{\beta}_{I_i|0, \boldsymbol{\nu}}} \right] (\boldsymbol{\theta}_{I_i}).$$

□

**Remark 2.** If  $A_\nu$  computed in Eqn. (10) is independent of  $\nu$ , then we can write each term of the sum in the following way:

$$f_i(\nu, \theta_{I_i}) := \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_\nu^{(i)}}{\mathcal{F}\tilde{\beta}_{I_i|0,\nu}} \right] (\theta_{I_i}) = u_i(\theta_{I_i}) + v_i(\nu).$$

*Proof.* With simplified notation, we have:  $A_\nu(\theta) = A(\theta) = f_0(\nu) + \sum_{i=1}^n f_i(\nu, \theta_{I_i})$ . Let  $\nu_0$  be some fixed value for the parameter  $\nu$ . Since  $A_\nu$  does not depend on  $\nu$ , then, for any  $\nu$  and  $\theta$ :

$$0 = A_\nu(\theta) - A_{\nu_0}(\theta) = \sum_{i=1}^n (f_i(\nu, \theta_{I_i}) - f_i(\nu_0, \theta_{I_i})).$$

Thus, each term  $f_i(\nu, \theta_{I_i}) - f_i(\nu_0, \theta_{I_i})$  is constant w.r.t.  $\theta_{I_i}$ , so there exists  $v_i$  s.t.:

$$f_i(\nu, \theta_{I_i}) = f_i(\nu_0, \theta_{I_i}) + v_i(\nu).$$

By setting  $u_i(\theta_{I_i}) = f_i(\nu_0, \theta_{I_i})$ , the proof is achieved.  $\square$

**Corollary 2.** With a penalty  $r$  as defined in Proposition 1, under Conditions (i)-(iv), there exists a unique solution  $\alpha \in \mathcal{T}(\mathbb{R}^N)$  of Eqn. (8):

$$\alpha(\theta) \propto \exp(-\text{Ent}(\beta_{0,\nu})) \prod_{i=1}^n \exp\left(-\mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_\nu^{(i)}}{\mathcal{F}\tilde{\beta}_{I_i|0,\nu}} \right] (\theta_{I_i})\right) = \prod_{i=1}^n \alpha_i(\theta_{I_i}),$$

by using Remark 2, where each  $\alpha_i$  is a probability distribution over  $\mathbb{R}^{I_i}$ . The resulting  $\alpha$  does not depend on  $\nu$  (Condition (iii)).

Under some conditions, Corollary 2 shows that, if a penalty  $r_\nu(\mu) = \sum_{i=1}^n r_\nu^{(i)}(\mu_{I_i})$  can be expressed as the influence of a Bayesian prior, then this prior takes the form  $\alpha(\theta) = \prod_{i=1}^n \alpha_i(\theta_{I_i})$ . To summarize, a sum of penalties over different components can be translated into a product of independent prior distributions over the same components:

$$\text{penalty } r_\nu(\mu) = \sum_{i=1}^n r_\nu^{(i)}(\mu_{I_i}) \quad \Rightarrow \quad \text{prior } \alpha(\theta) = \prod_{i=1}^n \alpha_i(\theta_{I_i}). \quad (11)$$

**Remark 3.** Let  $\beta_{\mu,\nu}(\theta_{I_1}, \dots, \theta_{I_n})$  be some probability distribution. For a distribution  $\alpha(\theta) = \alpha_1(\theta_{I_1}) \cdots \alpha_n(\theta_{I_n})$ , we have:

$$\begin{aligned} \text{KL}(\beta_{\mu,\nu} \parallel \alpha) &= -\text{Ent}(\beta_{\mu,\nu}) - \sum_{i=1}^n \int_{\mathbb{R}^{I_i}} \ln(\alpha_i(\theta_{I_i})) \beta_{I_i|\mu,\nu}(\theta_{I_i}) \, d\theta_{I_i} \\ &= \text{KL}(\beta_{|\mu,\nu} \parallel \alpha) + \text{Ent}(\beta_{|0,\nu}) - \text{Ent}(\beta_{0,\nu}), \end{aligned}$$

where  $\beta_{|\mu,\nu}(\theta) = \beta_{I_1|\mu,\nu}(\theta_{I_1}) \cdots \beta_{I_n|\mu,\nu}(\theta_{I_n})$  is the independent counterpart of  $\beta_{\mu,\nu}(\theta)$ . We have also used the fact that the entropy is independent of the mean of the distribution: for all  $\mu$ ,  $\text{Ent}(\beta_{|\mu,\nu}) = \text{Ent}(\beta_{|0,\nu})$  and  $\text{Ent}(\beta_{\mu,\nu}) = \text{Ent}(\beta_{0,\nu})$ .

**Example 3.** We consider the family of bivariate Gaussian distributions  $(\beta_{\boldsymbol{\mu}, \boldsymbol{\Sigma}})_{\boldsymbol{\mu}, \boldsymbol{\Sigma}}$ , with mean  $\boldsymbol{\mu} \in \mathbb{R}^2$  and covariance matrix  $\boldsymbol{\Sigma} \in \mathbb{R}^{2 \times 2}$ . We impose a penalty:

$$r_{\boldsymbol{\Sigma}}(\boldsymbol{\mu}) := a(\boldsymbol{\Sigma}) + \frac{b_1(\boldsymbol{\Sigma})}{2} \mu_1^2 + \frac{b_2(\boldsymbol{\Sigma})}{2} \mu_2^2,$$

where  $a$ ,  $b_1$  and  $b_2$  are functions. According to Proposition 1, we have, with  $s(x) = x^2$ :

$$A_{\boldsymbol{\Sigma}}(\boldsymbol{\theta}) = -\text{Ent}(\beta_{0, \boldsymbol{\Sigma}}) - a(\boldsymbol{\Sigma}) - \frac{b_1(\boldsymbol{\Sigma})}{2} \mathcal{F}^{-1} \left[ \frac{\mathcal{F}s}{\mathcal{F}\tilde{\beta}_{1|0, \boldsymbol{\Sigma}}} \right] (\theta_1) - \frac{b_2(\boldsymbol{\Sigma})}{2} \mathcal{F}^{-1} \left[ \frac{\mathcal{F}s}{\mathcal{F}\tilde{\beta}_{2|0, \boldsymbol{\Sigma}}} \right] (\theta_2).$$

We compute  $\frac{\mathcal{F}s}{\mathcal{F}\tilde{\beta}_{1|0, \boldsymbol{\Sigma}}}$ . Let  $\varphi \in C_c^\infty(\mathbb{R}^N)$  be a test function. We have:

$$\begin{aligned} \left\langle \frac{\mathcal{F}s}{\mathcal{F}\tilde{\beta}_{1|0, \boldsymbol{\Sigma}}}, \varphi \right\rangle &= 2\pi \left\langle -\delta'', \varphi(t) \exp\left(\frac{\sigma_1^2 t^2}{2}\right) \right\rangle_t = 2\pi (-\varphi''(0) - \sigma_1^2 \varphi(0)) \\ &= 2\pi \langle -\delta'' - \sigma_1^2 \delta, \varphi \rangle, \end{aligned}$$

where  $\sigma_i^2$  is the variance of the  $i$ -th marginal,  $\theta_i \sim \beta_{i|0, \boldsymbol{\Sigma}}$ .

Since  $\mathcal{F}^{-1}[-2\pi\delta''] = s$  and  $\mathcal{F}^{-1}[2\pi\delta] = 1$ , we have:

$$\begin{aligned} A_{\boldsymbol{\Sigma}}(\boldsymbol{\theta}) &= -\text{Ent}(\beta_{0, \boldsymbol{\Sigma}}) - a(\boldsymbol{\Sigma}) - \frac{b_1(\boldsymbol{\Sigma})}{2} (\theta_1^2 - \sigma_1^2) - \frac{b_2(\boldsymbol{\Sigma})}{2} (\theta_2^2 - \sigma_2^2) \\ &= \left[ -\text{Ent}(\beta_{0, \boldsymbol{\Sigma}}) - a(\boldsymbol{\Sigma}) + \frac{b_1(\boldsymbol{\Sigma})}{2} \sigma_1^2 + \frac{b_2(\boldsymbol{\Sigma})}{2} \sigma_2^2 \right] - \frac{b_1(\boldsymbol{\Sigma})}{2} \theta_1^2 - \frac{b_2(\boldsymbol{\Sigma})}{2} \theta_2^2. \end{aligned}$$

According to Condition (iii), the function  $A_{\boldsymbol{\Sigma}}$  should not depend on  $\boldsymbol{\Sigma}$ , that is:

$$b_1(\boldsymbol{\Sigma}) = b_1^0, \quad b_2(\boldsymbol{\Sigma}) = b_2^0, \quad a(\boldsymbol{\Sigma}) = a^0 - \text{Ent}(\beta_{0, \boldsymbol{\Sigma}}) + \frac{b_1^0}{2} \sigma_1^2 + \frac{b_2^0}{2} \sigma_2^2,$$

where  $a^0$ ,  $b_1^0$  and  $b_2^0$  are constants. Moreover, Condition (iv) imposes  $b_1^0 > 0$  and  $b_2^0 > 0$ . Also,  $\sigma_i^2 = \Sigma_{ii}$ . So, if we denote by  $|\boldsymbol{\Sigma}|$  the determinant of  $\boldsymbol{\Sigma}$ , we have:

$$r_{\boldsymbol{\Sigma}}(\boldsymbol{\mu}) = a^0 - \frac{1}{2} \ln((2\pi e)^2 |\boldsymbol{\Sigma}|) + \frac{b_1^0}{2} [\Sigma_{11} + \mu_1^2] + \frac{b_2^0}{2} [\Sigma_{22} + \mu_2^2],$$

which is, up to a constant, the KL between  $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  and  $\alpha = \mathcal{N}(0, 1/b_1^0) \otimes \mathcal{N}(0, 1/b_2^0)$ .

According to Example 3, if the total penalty can be split into two penalties, each one over a different parameter, then the prior  $\alpha$  is a product of independent univariate distributions. Notably, this phenomenon occurs even if the variational family contains multivariate distributions.

## 5 Examples

In this section, we apply Theorem 1 and Corollary 1 to a series of examples of variational families and penalties. Our goal is to show how the initial choice of variational family and penalty leads to a specific prior. Moreover, we show that enforcing a Bayesian interpretation of the penalty constrains the penalty itself: some penalties do not correspond to any Bayesian prior.

In the following sections, we study several pairs of family/penalty. Stated informally, the main results are:

- if a  $\mathcal{L}^2$  penalty over  $\boldsymbol{\mu}$  corresponds to the influence of a Bayesian prior  $\alpha$ , then  $\alpha$  is necessarily Gaussian (Sections 5.1, 5.3, 5.4);
- as a result, in that case, given a variational family  $(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})_{\boldsymbol{\mu}, \boldsymbol{\nu}}$ , the penalty consists in a competition between  $\text{Ent}(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})$ , encouraged to be larger, and  $\text{Cov}(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})$ , encouraged to be smaller, balanced by the covariance of the Gaussian prior  $\alpha$ ;
- when the penalty has a shape that is not fully determined, e.g.,  $\beta_{\boldsymbol{\mu}, \sigma} = \mathcal{N}(\boldsymbol{\mu}, \sigma^2)$  and  $r_\sigma(\boldsymbol{\mu}) = a(\sigma) - b(\sigma) \cos(\frac{\boldsymbol{\mu}}{\tau})$ , with  $a$  and  $b$  undetermined functions, we are able to propose strong constraints on  $a$  and  $b$  so that  $r$  can be interpreted as the influence of a Bayesian prior. In this specific example, we obtain:  $a(\sigma) = a_0 - \frac{1}{2} \ln(2\pi e \sigma^2)$ , where  $a_0$  is a constant, and  $b(\sigma) \propto \exp(-\frac{\sigma^2}{2\tau^2})$  (Section 5.5);
- the  $\mathcal{L}^1$  penalty cannot be interpreted as the influence of a Bayesian prior if the variational family belongs to  $\mathcal{S}(\mathbb{R})$ , which includes the Gaussian family (see Section 5.6).

We include several proofs to illustrate how Theorem 1 and Corollary 1 can be applied.

**Framework.** Starting from here, we distinguish two kinds of variational families: families of multivariate distributions (M) and families of product of independent univariate distributions (PIU). For instance, if one looks for the posterior distribution of a vector of parameters  $\boldsymbol{\theta} \in \mathbb{R}^N$ , within the family of Gaussian distributions, one may consider two kinds of families:

$$\begin{aligned} \text{Multivariate:} & \quad \boldsymbol{\theta} \sim \beta_{\boldsymbol{\mu}, \boldsymbol{\Sigma}} = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), & \text{(M)} \\ \text{Prod. of Indep. Univariate:} & \quad \boldsymbol{\theta} \sim \beta_{\boldsymbol{\mu}, \boldsymbol{\sigma}} = \mathcal{N}(\mu_1, \sigma_1^2) \otimes \cdots \otimes \mathcal{N}(\mu_N, \sigma_N^2), & \text{(PIU)} \end{aligned}$$

where  $\boldsymbol{\mu} = (\mu_1, \dots, \mu_N)$ ,  $\boldsymbol{\Sigma} \in \mathbb{R}^{N \times N}$ , and  $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_N)$ .

First, families M and PIU are very different from a practical point of view: in case M, the number of variational parameters is  $N(N+1)$ , while it is only  $2N$  in case PIU. So, it is easier and quicker to optimize  $(\boldsymbol{\mu}, \boldsymbol{\sigma})$  than  $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  when  $N$  is large, e.g., in neural networks. Indeed, it is far more common to use PIU when training neural networks with VI (Graves, 2011).

Second, even when using a  $M$  variational family  $(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})_{\boldsymbol{\mu}, \boldsymbol{\nu}}$  in combination with a penalty  $r_{\boldsymbol{\nu}}(\boldsymbol{\mu}) = r^{(0)}(\boldsymbol{\nu}) + \sum_{i=1}^N r_{\boldsymbol{\nu}}^{(i)}(\mu_i)$ , Corollary 2 ensures that the corresponding prior, if it exists, takes the form:  $\alpha(\boldsymbol{\theta}) = \alpha_1(\theta_1) \cdots \alpha_N(\theta_N)$ . Thus, we can use Corollary 2 and Remark 3: first, we look for the  $(\alpha_i)_i$  solving:

$$r_{\boldsymbol{\nu}}^{(i)}(\mu_i) = \text{KL}(\beta_{i|\boldsymbol{\mu}, \boldsymbol{\nu}} \|\alpha_i),$$

then we compute  $\alpha(\boldsymbol{\theta}) = \alpha_1(\theta_1) \cdots \alpha_N(\theta_N)$ . Finally, we recover the penalty:

$$r_{\boldsymbol{\nu}}(\boldsymbol{\mu}) = \text{KL}(\beta_{|\boldsymbol{\mu}, \boldsymbol{\nu}} \|\alpha) + \text{Ent}(\beta_{|0, \boldsymbol{\nu}}) - \text{Ent}(\beta_{0, \boldsymbol{\nu}}) = \text{KL}(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}} \|\alpha).$$

Therefore, we focus mainly on cases where both priors and candidate posteriors are assumed to be **PIU**: they cover a wide range of practical cases, and also cases with  $M$  candidate posteriors and **PIU** priors, which reduce to **PIU** priors and posteriors.

## 5.1 Univariate Gaussian with $\mathcal{L}^2$ Penalty

Let us study the family of univariate Gaussian distributions (Example 1) with an  $\mathcal{L}^2$  penalty. Formally:

$$\beta_{\mu, \sigma}(\theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\theta - \mu)^2}{2\sigma^2}\right), \quad r_{\sigma}(\mu) = a(\sigma) + \frac{1}{2}b(\sigma)\mu^2,$$

where  $a$  and  $b$  are two functions.

**Corollary 3.** *If the penalty  $r_{\sigma}$  above corresponds to a prior  $\alpha$ , then  $\alpha$  is a Gaussian  $\mathcal{N}(0, \sigma_0^2)$ . Moreover,  $a$  and  $b$  must have the following form:*

$$a(\sigma) = \frac{1}{2} \left[ \frac{\sigma^2}{\sigma_0^2} + \ln\left(\frac{\sigma_0^2}{\sigma^2}\right) - 1 \right], \quad b(\sigma) = \frac{1}{\sigma_0^2}.$$

*Proof.* This result is an application of Corollary 1. The conditions (a)-(b) and (i)-(ii) of Theorem 1 are fulfilled. Thus, by using (5):

$$A_{\sigma}(\theta) = -\frac{1}{2} \ln(2\pi e\sigma^2) - \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_{\sigma}}{\mathcal{F}\check{\beta}_{0, \sigma}} \right].$$

Let  $s(\theta) = \theta^2$ . We have, for any test function  $\varphi \in C_c^{\infty}(\mathbb{R})$ :

$$\begin{aligned} \left\langle \frac{\mathcal{F}s}{\mathcal{F}\check{\beta}_{0, \sigma}}, \varphi \right\rangle &= -2\pi \left\langle \frac{\delta''}{\exp\left(-\frac{\sigma^2 t^2}{2}\right)}, \varphi(t) \right\rangle_t = -2\pi \left\langle \delta'', \varphi(t) \exp\left(\frac{\sigma^2 t^2}{2}\right) \right\rangle_t \\ &= -2\pi(\varphi''(0) + \sigma^2 \varphi(0)) = -2\pi \langle \delta'' + \sigma^2 \delta, \varphi \rangle. \end{aligned}$$

Thus:  $A_{\sigma}(\theta) = -\frac{1}{2} \ln(2\pi e\sigma^2) - a(\sigma) - \frac{1}{2}b(\sigma)(\theta^2 - \sigma^2)$ , which is a constant function of  $\theta$  (Condition (iii)) if, and only if, there exists two constants  $a_0$  and  $b_0$  such that:

$$a(\sigma) = -\frac{1}{2} \ln(2\pi e\sigma^2) + \frac{1}{2}b(\sigma)\sigma^2 - a_0, \quad b(\sigma) = b_0.$$

So,  $\alpha = \exp(A)$  is a density of probability (Condition (iv)) if  $b_0 > 0$  and  $\int \alpha = 1$ . If so,  $\alpha$  is the density of a Gaussian  $\mathcal{N}(0, \sigma_0^2)$ , with  $b_0 = \frac{1}{\sigma_0^2}$  and  $a_0 = -\frac{1}{2} \ln(2\pi\sigma_0^2)$ .  $\square$

Thus, the assumption that a penalty  $r$  arises from a variational interpretation is a strong constraint over  $r$ . Here, the penalty  $r$  was initially parameterized by a pair of real functions  $(a, b)$ , and is finally parameterized by a single number  $\sigma_0^2$ .

## 5.2 Dirac Family and the MAP

Another basic example is to use Dirac distributions as variational posteriors (Example 2):  $\beta_{\mu, \nu} = \delta_\mu$ . Since  $\delta_0 \notin \mathcal{S}(\mathbb{R})$ , Condition (b) is not satisfied. However, it is possible to apply Formula (5) and check that the resulting prior  $\alpha$  is consistent with a chosen penalty  $r$ . Applying Formula (5) yields

$$A = -\text{Ent}(\delta_0) \mathbb{1} - \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r}{\mathcal{F}\delta_0} \right] = -\mathcal{F}^{-1} \left[ \frac{\mathcal{F}r}{\mathbb{1}} \right] = -r.$$

Thus, if  $\exp(-r)$  integrates to  $0 < \kappa < \infty$ , then we can define  $\alpha = \frac{1}{\kappa} \exp(-r)$ . Then, we can check that indeed  $\text{KL}(\delta_\mu \| \alpha) = r(\mu)$  up to a constant:

$$\begin{aligned} \text{KL}(\delta_\mu \| \alpha) &= -\text{Ent}(\delta_0) - \langle \delta_\mu, \ln(\alpha) \rangle \\ &= -\ln \alpha(\mu) - \text{Ent}(\delta_0) \\ &= r(\mu) + \ln \kappa - \text{Ent}(\delta_0), \end{aligned}$$

which confirms that the proposed prior  $\alpha$  is consistent with the penalty  $r(\cdot)$ .

Thus, this formula recovers via variational inference the well-known penalty-prior equivalence in the MAP approximation,  $\alpha_{\text{VI}}(\theta) \propto \exp(-r(\theta)) \propto \alpha_{\text{MAP}}(\theta)$  (MacKay, 1992b).

However, this is somewhat formal: the entropy  $\text{Ent}(\delta_0)$  of a Dirac function is technically undefined and is an “infinite constant”. In practice, though, with a finite machine precision  $\varepsilon$ , a Dirac mass can be defined as a uniform distribution over an interval of size  $\varepsilon$ , and  $\text{Ent}(\delta_0)$  becomes the finite constant  $\ln \varepsilon$ .

## 5.3 Arbitrary Variational Family with $\mathcal{L}^2$ Penalty

Let us consider a family of univariate variational posteriors  $(\beta_{\mu, \nu})_{\mu, \nu}$  and a penalty:

$$r_\nu(\mu) := a(\nu) + \frac{1}{2} b(\nu) \mu^2, \tag{12}$$

where  $a$  and  $b$  are functions. We assume that all Conditions (a)-(c) and (i)-(iv) are fulfilled.

**Corollary 4.** *If  $r$  corresponds to a KL-divergence with prior  $\alpha$ , then there exists  $\sigma_0 > 0$  s.t.:*

$$\begin{aligned} r_{\nu}(\mu) &= \frac{\text{Var}(\beta_{0,\nu}) + \mu^2}{2\sigma_0^2} - \text{Ent}(\beta_{0,\nu}) + \frac{1}{2} \ln(2\pi e\sigma_0^2), \\ \alpha(\theta) &= \frac{1}{\sqrt{2\pi\sigma_0^2}} \exp\left(-\frac{\mu^2}{2\sigma_0^2}\right). \end{aligned} \quad (13)$$

*Proof.* Formula (5) yields:

$$A_{\nu}(\theta) = -\text{Ent}(\beta_{0,\nu}) - \mathcal{F}^{-1} \left[ \frac{\mathcal{F}r_{\nu}}{\mathcal{F}\check{\beta}_{0,\nu}} \right].$$

Let  $s(\theta) = \theta^2$ . We have, for any test function  $\varphi \in C_c^\infty(\mathbb{R})$ :

$$\left\langle \frac{\mathcal{F}s}{\mathcal{F}\check{\beta}_{0,\nu}}, \varphi \right\rangle = -2\pi \left\langle \delta'', \frac{\varphi(t)}{\check{\beta}_{0,\nu}(t)} \right\rangle_t = -2\pi[\varphi(0) + \text{Var}(\check{\beta}_{0,\nu})\varphi''(0)],$$

since  $(\mathcal{F}\check{\beta}_{0,\nu})(0) = 1$ ,  $(\mathcal{F}\check{\beta}_{0,\nu})'(0) = 0$  ( $\check{\beta}_{0,\nu}$  is centered) and  $(\mathcal{F}\check{\beta}_{0,\nu})''(0) = \text{Var}(\check{\beta}_{0,\nu})$ .

Thus:  $A_{\nu}(\theta) = -\text{Ent}(\beta_{0,\nu}) - a(\nu) - \frac{1}{2}b(\nu)(\theta^2 - \text{Var}(\beta_{0,\nu}))$ , which is a constant function of  $\theta$  (Condition (iii)) if, and only if, there exists two constants  $a_0$  and  $b_0$  such that:

$$a(\nu) = -\text{Ent}(\beta_{0,\nu}) + \frac{1}{2}b(\nu)\text{Var}(\beta_{0,\nu}) - a_0, \quad b(\nu) = b_0.$$

So,  $\alpha = \exp(A)$  is a density of probability (Condition (iv)) if  $b_0 > 0$  and  $\int \alpha = 1$ . If so,  $\alpha$  is the density of a Gaussian  $\mathcal{N}(0, \sigma_0^2)$ , with  $b_0 = \frac{1}{\sigma_0^2}$  and  $a_0 = -\frac{1}{2} \ln(2\pi\sigma_0^2)$ .  $\square$

So, in general, if we penalize  $\mu$  with a term  $\mu^2$  and we are in a Bayesian setup, then the corresponding prior is necessarily Gaussian.

**Remark 4.** *Up to a constant,  $r_{\nu}$  is:*

$$r_{\nu}(\mu) = \left[ \frac{1}{2} \frac{\text{Var}(\beta_{0,\nu})}{\text{Var}(\alpha)} - (\text{Ent}(\beta_{0,\nu}) - \text{Ent}(\alpha)) \right] + \frac{\mu^2}{2\sigma_0^2}.$$

*So, there is a competition between the entropy term, which encourages the distribution  $\beta_{\mu,\nu}$  to be flatter, and the variance term, which pushes its variance towards 0. These two terms are balanced by the effect of the prior  $\alpha$ , through its own variance  $\text{Var}(\alpha) = \sigma_0^2$ .*

In this setup, the shape of the prior is entirely determined by the penalty, which means that the choice of the variational family does not matter. So, this is a generalization of Corollary 3 obtained in Section 5.1 for a Gaussian variational family. This striking result is a consequence of our choice for the family of

penalties. Indeed, our initial penalty  $r$  (12) is parameterized by two functions  $a$  and  $b$ , and we have discovered that, if  $r$  corresponds to a KL-divergence, then  $a$  and  $b$  are constrained to a specific shape (see Eqn. (13)). Additionally,  $a$  and  $b$  are not direct functions of  $\nu$ , but rather functions of  $\text{Ent}(\beta_{0,\nu})$  and  $\text{Var}(\beta_{0,\nu})$ .

**Remark 5.** *Moreover, the Fourier transform of the penalty  $r(x) = a + bx^2$ , is a linear combination of  $\delta$  and  $\delta''$ , both with support  $\{0\}$ . Due to the relation between the  $k$ -th derivatives of  $\mathcal{F}\beta$  at 0 and the  $k$ -th moments of the distribution  $\beta$ , the variance of  $\beta$  arises naturally in the constrained penalty  $r$  (see Eqn. (13)). So, if we consider  $\mathcal{L}^{2p}$  penalties with  $p > 0$  integer,  $p$ -th moments should arise as well in the constrained penalty.*

## 5.4 Multivariate Gaussian with $\mathcal{L}^2$ Penalty

Let us study the family of multivariate Gaussian distributions with a  $\mathcal{L}^2$  penalty. Formally:

$$\beta_{\boldsymbol{\mu},\boldsymbol{\Sigma}}(\boldsymbol{\theta}) = \frac{1}{\sqrt{(2\pi)^N |\boldsymbol{\Sigma}|}} \exp\left(-\frac{(\boldsymbol{\theta} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\theta} - \boldsymbol{\mu})}{2}\right),$$

$$r_{\boldsymbol{\Sigma}}(\boldsymbol{\mu}) = a(\boldsymbol{\Sigma}) + \frac{1}{2} \boldsymbol{\mu}^T B(\boldsymbol{\Sigma}) \boldsymbol{\mu},$$

where  $a : \mathbb{R}^{N \times N} \rightarrow \mathbb{R}$  and  $B : \mathbb{R}^{N \times N} \rightarrow \mathbb{R}^{N \times N}$  are two functions.

**Corollary 5.** *If the penalty  $r_{\boldsymbol{\Sigma}}$  above corresponds to a prior  $\alpha$ , then  $\alpha$  is a Gaussian  $\mathcal{N}(0, \mathbf{B}_0^{-1})$ . Moreover,  $a$  and  $B$  must have the following form:*

$$a(\boldsymbol{\Sigma}) = \frac{1}{2} \left[ \text{Tr}(\mathbf{B}_0 \boldsymbol{\Sigma}) - \ln(|\mathbf{B}_0| |\boldsymbol{\Sigma}|) - N \right], \quad B(\boldsymbol{\Sigma}) = \mathbf{B}_0.$$

*Proof.* This result is an application of Corollary 1. Conditions (a)-(c) and (i)-(ii) of Theorem 1 are fulfilled. Thus, by using (5):

$$A_{\boldsymbol{\Sigma}}(\boldsymbol{\theta}) = -\frac{1}{2} \ln((2\pi e)^N |\boldsymbol{\Sigma}|) - \mathcal{F}^{-1} \left[ \frac{\mathcal{F} r_{\boldsymbol{\Sigma}}}{\mathcal{F} \check{\beta}_{0,\boldsymbol{\Sigma}}} \right].$$

Let  $s(\boldsymbol{\theta}) = \boldsymbol{\theta}^T B(\boldsymbol{\Sigma}) \boldsymbol{\theta}$ . We have, for any test function  $\varphi \in C_c^\infty(\mathbb{R})$ :

$$\begin{aligned} \left\langle \frac{\mathcal{F} s}{\mathcal{F} \check{\beta}_{0,\sigma}}, \varphi \right\rangle &= \left\langle \mathcal{F} \left[ \sum_{i,j=1}^N B(\boldsymbol{\Sigma})_{ij} \theta_i \theta_j \right], \varphi(\mathbf{t}) \exp\left(\frac{\mathbf{t}^T \boldsymbol{\Sigma} \mathbf{t}}{2}\right) \right\rangle_{\mathbf{t}} \\ &= -(2\pi)^N \sum_{i,j=1}^N B(\boldsymbol{\Sigma})_{ij} \left\langle \partial^{p_{ij}} \delta, \varphi(\mathbf{t}) \exp\left(\frac{\mathbf{t}^T \boldsymbol{\Sigma} \mathbf{t}}{2}\right) \right\rangle_{\mathbf{t}}, \end{aligned}$$

where  $\partial^{p_{ij}} \delta : \varphi \mapsto \frac{\partial^2 \varphi}{\partial t_i \partial t_j}(0)$ . So:

$$\begin{aligned} \left\langle \frac{\mathcal{F}s}{\mathcal{F}\tilde{\beta}_{0,\sigma}}, \varphi \right\rangle &= -(2\pi)^N \sum_{i,j=1}^N [B(\boldsymbol{\Sigma})_{ij} \Sigma_{ij} \langle \delta, \varphi \rangle + B(\boldsymbol{\Sigma})_{ij} \langle \partial^{p_{ij}} \delta, \varphi \rangle] \\ &= -(2\pi)^N \left\langle \text{Tr}(B(\boldsymbol{\Sigma})\boldsymbol{\Sigma})\delta + \sum_{i,j=1}^N B(\boldsymbol{\Sigma})_{ij} \partial^{p_{ij}} \delta, \varphi \right\rangle. \end{aligned}$$

Thus:

$$\begin{aligned} \mathcal{F}^{-1} \left[ \frac{\mathcal{F}s}{\mathcal{F}\tilde{\beta}_{0,\sigma}} \right] (\boldsymbol{\theta}) &= \boldsymbol{\theta}^T B(\boldsymbol{\Sigma}) \boldsymbol{\theta} - \text{Tr}(B(\boldsymbol{\Sigma})\boldsymbol{\Sigma}) \\ A_{\boldsymbol{\Sigma}}(\boldsymbol{\theta}) &= -\frac{1}{2} \ln((2\pi e)^N |\boldsymbol{\Sigma}|) - a(\boldsymbol{\Sigma}) + \frac{1}{2} \text{Tr}(B(\boldsymbol{\Sigma})\boldsymbol{\Sigma}) - \frac{1}{2} \boldsymbol{\theta}^T B(\boldsymbol{\Sigma}) \boldsymbol{\theta}, \end{aligned}$$

where  $A_{\boldsymbol{\Sigma}}(\boldsymbol{\theta})$  is a constant function of  $\boldsymbol{\theta}$  (Condition (iii)) if, and only if, there exists two constants  $a_0 \in \mathbb{R}$  and  $\mathbf{B}_0 \in \mathbb{R}^{N \times N}$  such that:

$$a(\boldsymbol{\Sigma}) = -\frac{1}{2} \ln((2\pi e)^N |\boldsymbol{\Sigma}|) - a_0 + \frac{1}{2} \text{Tr}(B(\boldsymbol{\Sigma})\boldsymbol{\Sigma}), \quad B(\boldsymbol{\Sigma}) = \mathbf{B}_0.$$

So,  $\alpha = \exp(A)$  is a density of probability (Condition (iv)) if  $\mathbf{B}_0$  is positive definite and  $\int \alpha = 1$ . If so,  $\alpha$  is the density of a Gaussian  $\mathcal{N}(0, \mathbf{B}_0^{-1})$ ,  $a_0 = -\frac{1}{2} \ln((2\pi)^N |\mathbf{B}_0^{-1}|)$ .  $\square$

**Remark 6.** *More generally, one can deal with any multivariate family  $(\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}})_{\boldsymbol{\mu}, \boldsymbol{\nu}}$ , and not only the Gaussian, if the conditions (a)-(c) and (i)-(iv) are fulfilled. Then, one can easily adapt the proof of Corollary 4. With  $r_{\boldsymbol{\nu}}(\boldsymbol{\mu}) = a(\boldsymbol{\nu}) + \frac{1}{2} \boldsymbol{\mu}^T B(\boldsymbol{\nu}) \boldsymbol{\mu}$ , we have:*

$$\begin{aligned} r_{\boldsymbol{\nu}}(\boldsymbol{\mu}) &= \frac{1}{2} \boldsymbol{\mu}^T \mathbf{B}_0 \boldsymbol{\mu} + \text{Tr}(\mathbf{B}_0 \text{Cov}(\beta_{0,\boldsymbol{\nu}})) + \frac{1}{2} \ln((2\pi e)^N |\mathbf{B}_0^{-1}|) - \text{Ent}(\beta_{0,\boldsymbol{\nu}}) \\ \alpha(\boldsymbol{\theta}) &= (2\pi)^{-N/2} |\mathbf{B}_0|^{1/2} \exp\left(-\frac{1}{2} \boldsymbol{\theta}^T \mathbf{B}_0 \boldsymbol{\theta}\right). \end{aligned}$$

*So, even in the multivariate case, under some technical conditions, the  $\mathcal{L}^2$  penalty corresponds always to a (multivariate) Gaussian prior. Again, the parameters  $\boldsymbol{\nu}$  of  $\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}}$  are constrained through a competition between the entropy of  $\beta_{\boldsymbol{\mu}, \boldsymbol{\nu}}$  (which should be large) and its covariance (which should be small). This competition is modulated by the covariance  $\mathbf{B}_0^{-1}$  of the prior  $\alpha$ .*

## 5.5 Univariate Gaussian with Cosine Penalty

Let  $(\beta_{\mu, \sigma})_{\mu, \sigma}$  be the family of univariate Gaussian distributions  $\beta_{\mu, \sigma} = \mathcal{N}(\mu, \sigma^2)$ . In a context of quantization, a periodic penalty, such as cosine<sup>2</sup>, can be used

<sup>2</sup>Naumov et al. (2018) use  $\sin^2$  instead of  $\cos$ .

(Naumov et al., 2018). Such a penalty would encourage the parameters of the model to remain close to values that can be encoded with few bits:

$$r(\beta_{\mu,\sigma^2}) := a(\sigma^2) - b(\sigma^2) \cos\left(\frac{\mu}{\tau}\right),$$

where  $\tau > 0$ . With such penalty, if  $b(\sigma^2) > 0$ ,  $\mu$  is pushed towards the points  $\{2k\tau\pi : k \in \mathbb{Z}\}$ . We show that, if this penalty corresponds to a prior (possibly improper), then the functions  $a$  and  $b$  must have a specific shape with properties we detail below.

**Corollary 6.** *Let us consider the Gaussian variational family with penalty:*

$$r(\beta_{\mu,\sigma^2}) := a(\sigma^2) - b(\sigma^2) \cos\left(\frac{\mu}{\tau}\right),$$

where  $\tau > 0$ . Thus, if  $r$  corresponds to a KL-divergence with improper prior  $\alpha$ , then there exist  $a_0, b_0 \in \mathbb{R}$  s.t.:

$$\begin{aligned} \alpha(\theta) &\propto \exp\left(b_0 \cos\left(\frac{\mu}{\tau}\right)\right), \\ r(\beta_{\mu,\sigma^2}) &= a_0 - \frac{\ln(2\pi e\sigma^2)}{2} - b_0 e^{-\frac{\sigma^2}{2\tau^2}} \cos\left(\frac{\mu}{\tau}\right). \end{aligned}$$

Let us assume that  $b_0 > 0$ . In this setup,  $\mu$  is pushed towards the points  $\{2k\tau\pi : k \in \mathbb{Z}\}$  with strength  $\exp(-\frac{\sigma^2}{2\tau^2})$ . So, if  $\sigma \gg \tau$ , then  $\mu$  is very weakly penalized. On the contrary, if  $\sigma \ll \tau$ ,  $\mu$  is penalized with strength  $b_0$ . In short, if the standard deviation  $\sigma$  of  $\beta_{\mu,\sigma^2}$  is much larger than the width  $\tau$  of the oscillations of the penalty, then  $\mu$  is not penalized.

This behavior is consistent with what we expect when we train  $(\mu, \sigma)$ : if  $\sigma$  is so large that the center of the distribution  $\beta_{\mu,\sigma}$  covers several oscillations of the penalty, then the samples  $\theta \sim \beta_{\mu,\sigma}$  would lead to values of  $\cos(\frac{\theta}{\tau})$  that are independent of  $\mu$ . So, there would be no point in optimizing  $\mu$ . On the contrary, if  $\sigma \ll \tau$ , then almost all samples  $\theta \sim \beta_{\mu,\sigma}$  would lead to  $\cos(\frac{\theta}{\tau}) \approx \cos(\frac{\mu}{\tau})$ . So, optimizing  $\mu$  would be meaningful to obtain better samples.

From the point of view of  $\sigma$ , if the penalty  $-\cos(\frac{\mu}{\tau})$  is large, then smaller  $\exp(-\frac{\sigma^2}{2\tau^2})$  are encouraged, which means larger  $\sigma$ . On the contrary, if  $-\cos(\frac{\mu}{\tau})$  is already negative, then  $\sigma$  is encouraged to be smaller. To summarize, if  $\mu$  is strongly penalized, then  $\sigma$  tends to increase, making the samples  $\theta \sim \beta_{\mu,\sigma}$  explore more, and the more  $\mu$  is close to a minimum, the more the samples exploit the region around  $\mu$ .

## 5.6 Smooth Variational Family with $\mathcal{L}^1$ Penalty

In Section 4, we have proven results if some restrictive conditions are fulfilled. Specifically, for a variational family  $(\beta_{\mu,\nu})_{\mu,\nu}$ , we need  $r_\nu \in \mathcal{S}'(\mathbb{R})$  in order to be able to compute  $r_\nu$  (see Theorem 1). Since  $r : x \mapsto |x|$  does not belong to  $\mathcal{S}'(\mathbb{R})$ , then we cannot apply Theorem 1 to the  $\mathcal{L}^1$  penalty, which is quite common. So, one could argue that the conditions on  $r$  are too restrictive.

However, according to Proposition 2 below, when the distributions  $\beta_{\mu,\nu}$  of the variational family are smooth enough, it is impossible to find a prior  $\alpha$  s.t.:

$$\exists K \in \mathbb{R} : \forall \mu \in \mathbb{R}, \quad |\mu| + K = \text{KL}(\beta_{\mu,\nu} \parallel \alpha).$$

**Proposition 2.** *If the variational family  $(\beta_{\mu,\nu})_{\mu,\nu}$  is included in  $\mathcal{S}(\mathbb{R})$  and the penalty  $r$  does not belong to  $C^\infty(\mathbb{R})$ , then it is impossible to find any prior  $\alpha \in \mathcal{T}(\mathbb{R})$  (even improper) such that:*

$$\exists K \in \mathbb{R} : \quad r_\nu(\mu) + K = \text{KL}(\beta_{\mu,\nu} \parallel \alpha).$$

Proposition 2 is an immediate consequence of a property of the convolution:

**Proposition 3.** *(Stein and Weiss, 2016, Chap. I, Theorem 3.13) Let  $A \in \mathcal{S}'(\mathbb{R})$  and  $\beta_{0,\nu} \in \mathcal{S}(\mathbb{R})$ . Thus,  $A * \beta_{0,\nu} \in C^\infty(\mathbb{R})$ .*

This negative result is corroborated by Kaban (2007), who pointed out that a Laplace prior is sparsifying in the MAP case, but not when computing the full Bayesian posterior. The smoothing effect present in Proposition 3 has also been noticed. So, the ‘‘prior’’  $\alpha(x) \propto \exp(-|x|)$  corresponding to the penalty  $r(x) = |x|$  in the MAP case is no longer valid, and cannot even be tuned to generate a KL proportional to  $|\mu|$ , when considering smooth candidate posteriors.

**Remark 7.** *This regularity condition on  $r$  is implicitly contained in Condition (ii):  $T := \mathcal{F}r_\nu / \mathcal{F}\check{\beta}_{0,\nu} \in \mathcal{S}'(\mathbb{R}^N)$ . In other words, if  $r_\nu$  is not regular enough compared to  $\check{\beta}_{0,\nu}$ ,  $T$  would ‘‘diverge too quickly to infinity’’, so  $T \notin \mathcal{S}'(\mathbb{R}^N)$ .*

## 6 Building a Heuristic for the Penalty Factor

We show how our main results can be used to provide a heuristic on the penalty factor and the variance of the prior in the context of neural networks. Specifically, this heuristic leads to a non-isotropic diagonal prior, which contrasts with commonly used priors in Bayesian neural networks. On a small set of experiments, this heuristic provides a penalty factor relatively close to the optimal one. Additionally, we discuss the relation between our results and the results obtained in the framework of the ‘‘cold posterior effect’’.

### 6.1 Tuning the Penalty Factor

Let  $\mathcal{B} = \{\beta_{\mathbf{u}}\}$  be a variational family. Given a custom penalty  $r$  and a *penalty factor*  $\lambda > 0$ , we aim to minimize the following loss on  $\mathcal{B}$ :

$$\mathcal{L}(\beta_{\mathbf{u}}) := -\mathbb{E}_{\theta \sim \beta_{\mathbf{u}}} \ln p_\theta(\mathcal{D}) + \lambda r(\mathbf{u}). \quad (14)$$

We first have to tune the hyperparameter  $\lambda$ , which balances the amount of regularization we want to set when training the model. Usually,  $\lambda$  is determined by performing cross-validation (Stone, 1974). In the specific case of neural

networks, cross-validation necessitates as many training runs as the number of different  $\lambda$  to be tested.

In this document, we provide a heuristic designed to output a reasonable value for  $\lambda$ . More specifically, this heuristic is made of three bricks:

1. interpret the penalty  $\lambda r(\mathbf{u})$  as the influence of a Bayesian prior  $\alpha_\lambda$ ;
2. assume that the prior  $\alpha_\lambda$  should have the same variance as a reasonable initialization distribution of the weights, e.g., the one proposed by [Glorot and Bengio \(2010\)](#), denoted by  $\alpha^*$ ;
3. we set  $\lambda$  in such a way that:  $\text{Var}(\alpha_\lambda) = \text{Var}(\alpha^*)$ .

Then, we test this heuristic with different neural networks.

## 6.2 Building the heuristic

We illustrate our method in the case of a Gaussian variational family and a  $\mathcal{L}^2$  penalty over the means:

$$\beta_{\boldsymbol{\mu}, \boldsymbol{\sigma}} := \beta_{\boldsymbol{\mu}, \boldsymbol{\sigma}} = \mathcal{N}(\mu_1, \sigma_1^2) \otimes \cdots \otimes \mathcal{N}(\mu_N, \sigma_N^2),$$

$$r_{\boldsymbol{\sigma}}(\boldsymbol{\mu}) := \sum_{i=1}^N r_{i, \boldsymbol{\sigma}}(\mu_i) = \sum_{i=1}^N [a_i(\boldsymbol{\sigma}) + b_i(\boldsymbol{\sigma})\mu_i^2].$$

**Step 1: interpretation of the penalty.** In this framework, the requirements of Corollary 2 are fulfilled, and the prior  $\alpha$  corresponding to  $r$  is:

$$\alpha(\boldsymbol{\theta}) := \alpha_1(\theta_1) \cdots \alpha_N(\theta_N),$$

where each  $\alpha_i$  is a Gaussian distribution with mean 0.

From now, and without loss of generality, we focus on one component  $\theta_i$  of  $\boldsymbol{\theta}$ , and we omit the index  $i$ :  $\theta := \theta_i$ ,  $\mu := \mu_i$ ,  $\sigma = \sigma_i$ ,  $\alpha := \alpha_i$ ,  $r_\sigma(\mu) := r_{i, \sigma_i}(\mu_i)$ . Also, we denote by  $\sigma_0^2$  the variance of  $\alpha_i$ . We have:

$$r_\sigma(\mu) = \text{KL}(\beta_{\mu, \sigma} \| \alpha) = \frac{1}{2} \ln \left( \frac{\sigma_0^2}{\sigma^2} \right) + \frac{\sigma^2 + \mu^2}{2\sigma_0^2} - \frac{1}{2}.$$

Thus we have:

$$\lambda r_\sigma(\mu) = \frac{1}{2} \ln \left( \frac{\frac{\sigma_0^2}{\lambda}}{\sigma^2} \right) + \frac{\sigma^2 + \mu^2}{2\frac{\sigma_0^2}{\lambda}} - \frac{1}{2} + K_\lambda,$$

where  $K_\lambda$  does not depend on  $(\mu, \sigma)$ , so its value does not affect optimization. Therefore:

$$\lambda r_\sigma(\mu) = \text{KL}(\beta_{\mu, \sigma} \| \alpha_\lambda), \quad \text{where} \quad \alpha_\lambda \sim \mathcal{N} \left( 0, \frac{\sigma_0^2}{\lambda} \right). \quad (15)$$

**Step 2: assumption on the variance of the prior.** Now that we have translated the penalty into a prior  $\alpha_\lambda$ , we are looking for a constraint to impose on  $\alpha_\lambda$ . To do so, we make the following *informal* postulate:

**Postulate 1.** *A distribution  $\alpha$  over the parameters  $\theta$  of a neural network is a reasonable prior over  $\theta$  if  $\alpha$  can be used as an initialization distribution of  $\theta$ .*

*Intuition.* We consider a neural network (NN) before training. In this situation, from a Bayesian point of view, no data point has been observed, so the Bayesian posterior over the parameters is equal to the prior.

When we start training the NN with a batch of data, we expect that the variational parameters  $\mathbf{u}$  would move to the optimum. Since we use the reparameterization trick (Kingma and Welling, 2013) to train  $\mathbf{u}$ , we expect that the gradient of the loss according to the parameters  $\theta$  of the NN is such that training is not stuck.

So, when we send the first batch of data to the NN for training, we sample  $\theta$  from the prior  $\alpha$ , and information should propagate and backpropagate properly across the network. This condition corresponds to the ones that are usually imposed on initialization distributions (Glorot and Bengio, 2010; He et al., 2015; Schoenholz et al., 2017).

Therefore, to ensure that training is not stuck at the beginning, it is reasonable to require the prior distribution to be usable as an initialization distribution for  $\theta$ .  $\square$

This postulate has previously been used by Ollivier (2018).

More specifically, for the distribution of the weight  $W$  of a neuron with  $p$  inputs and activation function  $\varphi$ :

- Glorot and Bengio (2010) recommend:  $\text{Var}(W) = \frac{1}{p}$  (under the hypothesis  $\varphi = \text{Id}$ );
- He et al. (2015) recommend:  $\text{Var}(W) = \frac{2}{p}$ , when  $\varphi = \text{ReLU} = \max(\cdot, 0)$ ;
- Paszke et al. (2019) recommend:  $\text{Var}(W) = \left(\frac{5}{3}\right)^2 \frac{1}{p}$ , when  $\varphi = \tanh$ .

Now, we are ready to compute a reasonable  $\lambda$ .

**Step 3: computation of  $\lambda$ .** Regardless of the activation function  $\varphi$ , the initialization distribution of a weight  $W$  of a neuron with  $p$  inputs has a variance:  $\sigma_*^2 = \frac{g^2}{p}$ , where  $g > 0$  is a “gain” depending on  $\varphi$ .

According to Postulate 1, the variance of the prior  $\alpha_\lambda$  should be equal to  $\sigma_*^2$ . So, according to Eqn. (15), we should have:

$$\frac{\sigma_0^2}{\lambda} = \frac{g^2}{p}. \quad (16)$$

Now, we just have to tune  $\sigma_0^2$  and  $\lambda$  so that Eqn. (16) holds. Our choice is:

$$\sigma_0 = \frac{1}{\sqrt{p}} \quad \text{and} \quad \lambda = \frac{1}{g^2}. \quad (17)$$

### 6.3 Experimental setup

**Loss and priors.** Above all, we average the loss (14) over the dataset  $\mathcal{D}$  of size  $n$ :

$$\bar{\mathcal{L}}(\beta_{\mathbf{u}}) := \underbrace{-\frac{1}{n}\mathbb{E}_{\theta\sim\beta_{\mathbf{u}}}\ln p_{\theta}(\mathcal{D})}_{\text{NLL term}} + \underbrace{\bar{\lambda}r(\mathbf{u})}_{\text{penalty term}},$$

where  $\bar{\lambda} = \frac{\lambda}{n}$  and ‘‘NLL’’ stands for ‘‘negative log-likelihood’’. Our heuristic (17) becomes:

$$\sigma_0 = \frac{1}{\sqrt{p}} \quad \text{and} \quad \bar{\lambda} = \frac{1}{ng^2}. \quad (18)$$

We minimize  $\bar{\mathcal{L}}$  according to the means and variances  $(\mu_i, \sigma_i^2)$  of each parameter  $\theta_i$ . The penalty associated with  $(\mu_i, \sigma_i^2)$  is equal to  $\frac{1}{n}\text{KL}(\beta_{\mu_i, \sigma_i} \parallel \alpha_i)$ , where  $\alpha_i \sim \mathcal{N}(0, \frac{1}{p_i})$  and  $p_i$  is the number of inputs of the neuron to which the weight  $\theta_i$  belongs.

**Datasets and architectures.** We consider three setups:

1. MNIST + multilayer perceptron (MLP): layers of sizes 200, 100, 10,  $\varphi = \text{ReLU}$ ;
2. CIFAR-10 + LeNet (LeCun et al., 1998),  $\varphi = \text{tanh}$ ;
3. CIFAR-10 + VGG13\*<sup>3</sup> (Simonyan and Zisserman, 2014),  $\varphi = \text{ReLU}$ .

In each setup, we train the variational parameters of the given architecture with Adam (Kingma and Ba, 2015) and with various values for  $\bar{\lambda}$ , centered in  $\bar{\lambda} = \frac{1}{n}$ . We extracted a validation set from the usual training set, of size 6000 for MNIST and of size 5000 for CIFAR-10. No data augmentation has been used.

**Training details:**

1. MNIST + MLP: Adam with learning rate  $\eta = 10^{-3}$ , 200 epochs;
2. CIFAR-10 + LeNet: Adam with learning rate  $\eta = 10^{-3}$ , 400 epochs;
3. CIFAR-10 + VGG13\*: Adam with learning rate  $\eta = 10^{-4}$ , 600 epochs.

### 6.4 Results

We have plotted in Figure 1 the results for the 3 different setups. Without surprise, we observe a transition between the overfitting regime (small  $\bar{\lambda}$ ) and the underfitting regime (large  $\bar{\lambda}$ ): when  $\bar{\lambda}$  is small, the penalty term is too small to regularize correctly the loss, and the model tends to overfit the training dataset; when  $\bar{\lambda}$  is large, the penalty term is too restrictive and prevents the model from fitting correctly the data.

<sup>3</sup>VGG13\* is VGG13 without batch-norm and with only one fully-connected layer.




Figure 1: Test NLL obtained at the epoch where the validation NLL is optimal, for various penalty factors  $\bar{\lambda}$ . The quality of our heuristics can be estimated by measuring the closeness of the minimum of the blue curve to the solid orange line.

**Correct order of magnitude.** Between these regimes, there is some “sweet spot” for  $\bar{\lambda}$ , which is typically the region where our heuristic should lie. Let us denote our heuristic by  $\hat{\lambda}$  and the optimal  $\bar{\lambda}$  by  $\lambda^*$ . In Figures 1a and 1c, our heuristic  $\hat{\lambda}$  overestimates  $\lambda^*$ , while it slightly underestimates it in Figure 1b. Overall, our estimation does not deviate from the optimum by a factor 10. So, its order of magnitude is roughly correct.

**Influence of the architecture.** Interestingly, the error between our estimation and  $\lambda^*$  is much greater in Figure 1c than in Figure 1b, while the  $\hat{\lambda}$  is approximately the same in both cases. Two factors explain this similarity between the  $\hat{\lambda}$ : the size of the training dataset is the same in both cases; the activation functions used in both cases have a similar squared gain  $g^2$ . However,  $\hat{\lambda}$  lies in the sweet spot in Figure 1b, and is clearly outside of it in Figure 1c.

Consequently, a good heuristic for  $\bar{\lambda}$  cannot rely only on the size of the training dataset and the choice of activation function. One can assume that the quality of our proposal  $\hat{\lambda}$  depends also on the architecture of the NN, including its number of parameters (LeNet has only 60 000 parameters, while VGG13\* has roughly 10 000 000 parameters). This empirical finding corroborates (Izmailov et al., 2021), which states that the prior design should take into account the architecture of the NN, via the function mapping the vector of parameters of a NN to the function represented by the NN.

**Link with the cold posterior effect.** According to Wenzel et al. (2020), the “cold posterior effect” is an empirical finding according to which, when training neural networks, a posterior tempered with a well-chosen temperature  $T < 1$  leads to a better test loss than the Bayesian (non-tempered,  $T = 1$ ) posterior. This finding is related to the choice of the penalty factor  $\bar{\lambda}$ : if  $T < 1$ , then  $\bar{\lambda} < 1$ .

However, as shown in Eqn. (16), the quantity of interest is not  $\bar{\lambda}$  alone, but the ratio  $\sigma_0^2/\lambda$ . So, focusing only on the optimal value of  $\bar{\lambda}$  is meaningless if the variance  $\sigma_0^2$  is chosen carelessly.

We note that many works about Bayesian inference applied to NNs, such as (Zhang et al., 2018; Bae et al., 2018; Ashukha et al., 2019), but also (Izmailov et al., 2021), which obtains posteriors of much better quality than with VI by using HMC and (Wenzel et al., 2020), assume that the posterior is isotropic. This assumption is out of step with our heuristic (see Eqn. (16) and (18)): if we want the prior distribution  $\alpha$  to be usable to initialize the NN, then  $\alpha$  *cannot be isotropic*.<sup>4</sup> The variance of the prior over a weight  $\theta_i$  depends on some architecture hyperparameters (number of neurons per layer). This choice of diagonal layer-dependent prior covariance has previously been proposed by (Ollivier, 2018).

Nonetheless, Figures 1a and 1c exhibit optimal values  $\lambda^*$  that are smaller than  $\hat{\lambda}$ , which seems to confirm the “cold posterior effect”. But, again, one could argue that our choice of prior is not entirely satisfying and propose a better heuristic for  $\alpha$ , which would consequently change the heuristic for  $\lambda$ , challenging again the “cold posterior effect”.

## 7 Discussion

**Conditions of Theorem 1.** Overall, the technical conditions of Theorem 1 do not play the same role. Condition (b) ensures that the KL-divergence in Eqn. (4) remains computable for a very wide range of log-priors  $A$ . Similarly, Condition (ii) ensures that  $A$  can be computed within the framework of the standard theory of distributions. Specifically, the Fourier transform of  $\mathcal{F}r_\nu/\mathcal{F}\tilde{\beta}_{0,\nu}$  can be computed. Finally, Condition (c) ensures the uniqueness of the solution of Eqn. (4) in  $\mathcal{S}'(\mathbb{R}^N)$ .

**Trade-off between the Condition (b) and the search space  $\mathcal{T}$ .** We have chosen a very strong Condition (b), that is, each  $\beta_{\mu,\nu}$  should belong to  $\mathcal{S}(\mathbb{R}^N)$ , in order to allow for log-priors  $A$  to belong to the very large space  $\mathcal{S}'(\mathbb{R}^N)$ . If we want to deal with variational families with less regular densities<sup>5</sup>, we would have to shrink the search space of  $A$ .

**Allowing improper priors.** In Section 5.5, the solution  $\alpha$  to Eqn. (8) is an improper prior. If one wants a proper prior, one can add an  $\mathcal{L}^2$  penalty to the initial cosine penalty, which would impose Gaussian tails. But, it is also possible to deal with improper priors in general, if we are certain that the resulting variational posterior is a probability distribution. To allow improper priors, one just needs to relax Corollary 1 by removing Condition (iv).

**Beyond translation-invariant variational families.** To prove Theorem 1, we have used the fact that the variational family  $(\beta_{\mu,\nu})_{\mu,\nu}$  is translation-

<sup>4</sup>Fortuin et al. (2022) have considered non-isotropic multivariate priors, but not non-isotropic priors with diagonal covariance, which would be close to what we are doing.

<sup>5</sup>For instance, derivable only a finite number of times and with a polynomial decrease at infinity.

invariant:  $\beta_{\mu,\nu}(\theta) = \beta_{0,\nu}(\theta - \mu)$  for all  $\mu, \nu, \theta$ . This condition allowed us to formulate a part of the KL as a convolution between two functions, which can be disentangled by using the Fourier transform:  $\mathcal{F}_\mu \int_{-\infty}^{\infty} A(\theta) \beta_{\mu,\nu}(\theta) d\theta = (\mathcal{F}A) \cdot (\mathcal{F}\tilde{\beta}_{0,\nu})$ .

But there are other ways to disentangle functions inside an integral, if we assume different properties for the variational family. For instance, one can deal with a *scale-invariant* variational family:  $\beta_{\mu,\nu}(\theta) = \frac{1}{\mu} \beta_{1,\nu}(\frac{\theta}{\mu})$  for all  $\mu, \nu, \theta$ , and with support included in  $\mathbb{R}^+$ . In this case, we can express the KL with a *multiplicative convolution*  $\star$ , which can be disentangled by using the *Mellin transform*  $\mathcal{M}$ : first, transform the cross-entropy term into a multiplicative convolution:  $\int_0^\infty A(\theta) \beta_{\mu,\nu}(\theta) d\theta = \frac{1}{\mu} \int_0^\infty \theta A(\theta) \tilde{\beta}_{1,\nu}(\mu/\theta) \frac{d\theta}{\theta} = \frac{1}{\mu} [(\text{Id} \cdot A) \star \tilde{\beta}_{1,\nu}](\mu)$ , where Id is the identity function and  $\tilde{\beta}(\theta) = \beta(1/\theta)$ ; second, apply the Mellin transform  $\mathcal{M}$ :  $\mathcal{M}_\mu[\mu \int_0^\infty A(\theta) \beta_{\mu,\nu}(\theta) d\theta] = (\mathcal{M}[\text{Id} \cdot A]) \cdot (\mathcal{M}\tilde{\beta}_{1,\nu})$ ; third, use basic properties of  $\mathcal{M}$ :  $(\mathcal{M} \int_0^\infty A(\theta) \beta_{\mu,\nu}(\theta) d\theta)(t+1) = (\mathcal{M}A)(t+1) \cdot (\mathcal{M}\beta_{1,\nu})(-t)$ . Finally, it is possible to isolate  $\mathcal{M}A$  and recover  $A$  by using  $\mathcal{M}^{-1}$ . Then, one can rewrite Theorem 1 and Corollary 1 with the appropriate formulas and hypotheses.

**Heuristics for the prior and the penalty factor in the case of neural networks.** In Section 6, we showed how to combine our main theoretical results with usual initialization heuristics in neural networks to discover a heuristic for the prior and the penalty factor. We showed the results of such heuristics in a series of experiments, and we proposed a discussion about the “cold posterior effect” (Wenzel et al., 2020) in variational inference. Notably, we suggest a specific choice of the prior distribution, which should be non-isotropic, and with a variance over the weights of each layer depending on the input dimension of the considered layer.

## 8 Conclusion

We have provided a theorem that bridges the gap between empirical penalties and Bayesian priors when learning the distribution of the parameters of a model through Variational Inference. This way, various regularization techniques can be studied in a single Bayesian framework, and be seen as probability distributions. This unification of points of view is valuable on the theoretical side, since it can be used to apply well-know theorems in variational inference, such as results about convergence rates. Also, we have shown experimentally material that such theoretical results can be used in practice to help tuning the hyperparameters, especially when the model exhibits some structure, such as neural networks.

## References

Alquier, P., Ridgway, J., et al. (2020). Concentration of tempered posteriors and of their variational approximations. *Annals of Statistics*, 48(3):1475–1497.

- Ashukha, A., Lyzhov, A., Molchanov, D., and Vetrov, D. (2019). Pitfalls of in-domain uncertainty estimation and ensembling in deep learning. In *International Conference on Learning Representations*.
- Bae, J., Zhang, G., and Grosse, R. (2018). Eigenvalue corrected noisy natural gradient. *arXiv preprint arXiv:1811.12565*.
- Chérif-Abdellatif, B.-E. (2020). Convergence rates of variational inference in sparse deep learning. In *International Conference on Machine Learning*, pages 1831–1842. PMLR.
- Fortuin, V., Garriga-Alonso, A., Ober, S. W., Wenzel, F., Ratsch, G., Turner, R. E., van der Wilk, M., and Aitchison, L. (2022). Bayesian neural network priors revisited. In *International Conference on Learning Representations*.
- Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, pages 249–256. JMLR Workshop and Conference Proceedings.
- Graves, A. (2011). Practical variational inference for neural networks. *Advances in Neural Information Processing Systems*, 24.
- Gribonval, R. (2011). Should penalized least squares regression be interpreted as maximum a posteriori estimation? *IEEE Transactions on Signal Processing*, 59(5):2405–2410.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034.
- Hinton, G. E. and van Camp, D. (1993). Keeping neural networks simple. In *International Conference on Artificial Neural Networks*, pages 11–18. Springer.
- Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851.
- Izmailov, P., Vikram, S., Hoffman, M. D., and Wilson, A. G. G. (2021). What are bayesian neural network posteriors really like? In *International conference on machine learning*, pages 4629–4640. PMLR.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. (1999). An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233.
- Kaban, A. (2007). On bayesian classification with Laplace priors. *Pattern Recognition Letters*, 28(10):1271–1282.
- Kingma, D. P. and Ba, J. L. (2015). Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.

- Kingma, D. P., Salimans, T., and Welling, M. (2015). Variational dropout and the local reparameterization trick. In *Advances in Neural Information Processing Systems*, pages 2575–2583.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Knoblauch, J., Jewson, J., and Damoulas, T. (2022). An optimization-centric view on Bayes’ rule: Reviewing and generalizing variational inference. *Journal of Machine Learning Research*, 23(132):1–109.
- Krogh, A. and Hertz, J. A. (1992). A simple weight decay can improve generalization. In *Advances in neural information processing systems*, pages 950–957.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.
- Louizos, C., Ullrich, K., and Welling, M. (2017). Bayesian compression for deep learning. In *Advances in Neural Information Processing Systems*, pages 3288–3298.
- MacKay, D. J. (1992a). Bayesian model comparison and backprop nets. In *Advances in Neural Information Processing Systems*, pages 839–846.
- MacKay, D. J. (1992b). A practical Bayesian framework for backpropagation networks. *Neural Computation*, 4(3):448–472.
- MacKay, D. J. (1995). Probable networks and plausible predictions—a review of practical Bayesian methods for supervised neural networks. *Network: Computation in Neural Systems*, 6(3):469–505.
- MacKay, D. J. (2003). *Information theory, inference and learning algorithms*. Cambridge University Press.
- Molchanov, D., Ashukha, A., and Vetrov, D. (2017). Variational dropout sparsifies deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2498–2507. JMLR. org.
- Naumov, M., Diril, U., Park, J., Ray, B., Jablonski, J., and Tulloch, A. (2018). On periodic functions as regularizers for quantization of neural networks. *arXiv preprint arXiv:1811.09862*.
- Ollivier, Y. (2018). Online natural gradient as a Kalman filter. *Electronic Journal of Statistics*, 12(2):2930–2961.
- Olshausen, B. A. and Field, D. J. (1997). Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, 37(23):3311–3325.

- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Polyanin, A. D. and Manzhirov, A. V. (1998). *Handbook of integral equations*. CRC press.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.
- Schoenholz, S. S., Gilmer, J., Ganguli, S., and Sohl-Dickstein, J. (2017). Deep information propagation. In *International Conference on Learning Representations*.
- Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding machine learning: From theory to algorithms*. Cambridge University Press.
- Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Stein, E. M. and Weiss, G. (2016). *Introduction to Fourier analysis on Euclidean spaces*, volume 32. Princeton University Press.
- Stone, M. (1974). Cross-validatory choice and assessment of statistical predictions. *Journal of the royal statistical society: Series B (Methodological)*, 36(2):111–133.
- Wenzel, F., Roth, K., Veeling, B., Swiatkowski, J., Tran, L., Mandt, S., Snoek, J., Salimans, T., Jenatton, R., and Nowozin, S. (2020). How good is the Bayes posterior in deep neural networks really? In *International Conference on Machine Learning*, pages 10248–10259. PMLR.
- Zhang, F. and Gao, C. (2020). Convergence rates of variational posterior distributions. *Annals of Statistics*, 48(4):2180–2207.
- Zhang, G., Sun, S., Duvenaud, D., and Grosse, R. (2018). Noisy natural gradient as variational inference. In *International conference on machine learning*, pages 5852–5861. PMLR.